











CDMTCS Research Report Series Design by Example for SQL Table Definitions with Functional Dependencies

Markus Kirchberg

School of Computing, The National University of Singapore, Singapore

Sven Hartmann Department of Informatics, Clausthal University of Technology, Clausthal-Zellerfeld, Germany

Sebastian Link Department of Computer Science, University of Auckland, Auckland, New Zealand



CDMTCS-426 October 2012 Centre for Discrete Mathematics and Theoretical Computer Science

Design by Example for SQL Table Definitions with Functional Dependencies

MARKUS KIRCHBERG School of Computing The National University of Singapore Singapore markus@comp.nus.edu.sg

SVEN HARTMANN Department of Informatics Clausthal University of Technology Clausthal-Zellerfeld, Germany sven.hartmann@tu-clausthal.de

SEBASTIAN LINK Department of Computer Science The University of Auckland, Private Bag 92019 Auckland, New Zealand s.link@auckland.ac.nz

October 10, 2012

Abstract

A database is C-Armstrong for a given set of constraints in a class C if it satisfies every constraint of the set and violates every constraint in C not implied by the set. Therefore, Armstrong databases are test data that perfectly illustrate the current perceptions about the semantics of a schema. We extend the existing theory of Armstrong relations to a toolbox of Armstrong tables. That is, we investigate structural and computational properties of Armstrong tables for the class of functional dependencies (FDs) over SQL tables. Relations are special instances of SQL tables with no duplicate rows and no null value occurrences. While FDs do not enjoy Armstrong tables, the combined class of standard FDs and NOT NULL constraints does enjoy Armstrong tables. The problem of finding an Armstrong table is shown to be precisely exponential for this combined class. However, we establish an algorithm that computes Armstrong tables with a size at most quadratic in that of a minimum-sized Armstrong table. Our resulting toolbox of Armstrong tables can be applied by data engineers to concisely visualize constraints on SQL data. Such support can lead to designs that guarantee efficient data management in practice.

Keywords: Armstrong table, Functional dependency, Schema design, SQL, Test data, Uniqueness

1 Introduction

A database system is a software package that manages a collection of persistent information in a shared, reliable, effective and efficient way. The core of most database systems is still founded on *the relational model of data* [1]. In this model, data is stored in a collection of relations that may vary over time. A relation is a set of tuples over a given time-invariant relation schema.

Relations permit the storage of inconsistent data, i.e., data that violate conditions which every meaningful relation ought to satisfy. Consequently, data engineers must identify the conditions, called data dependencies, that restrict the relations to those that are meaningful for the application domain at hand. Functional dependencies (FDs) constitute the probably most important class of data dependencies [2, 3, 4]. According to [5] they make up around two-thirds of all uni-relational data dependencies (those defined over a single relation schema) in practice. FDs form one of the most prolific database concepts. They are essential for the fundamental tasks of database modeling and design [6, 7, 8, 9, 10], normalization on the conceptual [11, 12, 13] and logical level [11, 14, 15], and maintenance [16, 17, 18]. They have proven applications in data cleaning [19, 20], entry [7], exchange [21], integration [22, 23, 24], sampling [25, 26], warehousing [27], indexing [28], information retrieval [29], uncertain data management [30, 31], repairing [32] and consistent query answering [33], query optimization [34, 35, 36], access control [37] and privacy-preserving data publishing [38] among others. A basic prerequisite for addressing these fundamental tasks and taking effective advantage of these applications is that data engineers are able to identify the FDs that are meaningful for their application domain.

Example 1 Suppose that in designing an information system for a company the team of data engineers has identified the relationship between employees and managers in departments as an important information unit. Therefore, the team has decided to utilize the relation schema EMPLOYMENT with attributes Emp, Dept, and Mgr. The schema stores the name of an employee in the Emp column, the name of the department in which the employee works in the Dept column, and the name of the department's manager in the Mgr column. The team has identified additional constraints they consider important for the schema EMPLOYMENT. Firstly, the information stored in the columns Emp and Mgr must be total, i.e., no occurrences of null values are permitted in these columns. Consequently, an SQL definition of the table may look as follows:

CREATE TABLE EMPLOYMENT (Emp VARCHAR NOT NULL, Dept VARCHAR, Mgr VARCHAR NOT NULL); Moreover, the team of data engineers believes that every employee can work for only one department, and that every department has only one manager. These constraints can be expressed as the FDs $Emp \rightarrow Dept$ and $Dept \rightarrow Mgr$, respectively. The team would like to consult the experts of the application domain to find out whether their current design choice captures all the requirements necessary. In order to validate their own understanding of the application domain and to facilitate the knowledge acquisition from the domain experts the team plans to create test data that represents concisely their current understanding of the database.

Example 1 illustrates the findings of previous research [39, 40, 25, 41, 26, 42]: The creation of good test data can help data engineers with the challenge to consolidate the set of meaningful FDs. In fact, the findings suggest to use good test data to identify actually meaningful FDs that are incorrectly perceived as meaningless [40, 25, 26, 42]. It is thus an important task to create good test data effectively and efficiently. Informally, good means here that the test data satisfies the FDs that the team of data engineers considers meaningful, and violates the FDs the team perceives as meaningless.

More formally, given a set $\Sigma \cup \{\varphi\}$ of FDs we say that Σ *implies* φ if every relation that satisfies all FDs in Σ also satisfies φ . We use Σ^* to denote the *semantic closure* of Σ , i.e., the set of all FDs implied by Σ . Armstrong relations have been identified as an effective means to represent concisely abstract sets of FDs in the form of a single relation. Indeed, Armstrong relations are widely regarded as good test data and a helpful tool for data engineers to judge, justify, convey, or test their understanding of the relation schema [39, 40, 25, 41, 26, 42]. More precisely, an Armstrong relation for a set Σ of FDs is a single relation that satisfies every FD in Σ^* and violates every FD not in Σ^* [43]. Hence, an FD is implicit in the explicit specification of an FD set Σ if and only if it is satisfied by an Armstrong relation for Σ . For this reason, Armstrong relations represent one of the few instances where example-based reasoning is effective. That is, if Σ constitutes the design choice of the set of FDs currently perceived as meaningful to the underlying application domain, then an Armstrong relation r_{Σ} for Σ constitutes an example on which data engineers can test the meaningfulness of an arbitrary FD φ . Namely, φ is meaningful under the current design choice if and only if the example relation r_{Σ} satisfies φ . Consequently, this approach to design is called *design-by-example*. Empirical evidence has been presented that Armstrong relations help data engineers to identify meaningful FDs they perceived incorrectly as meaningless before inspecting an Armstrong relation [44]. Industry-leading data modeling tools, such as the *ERwin* data modeler, emphasize the need for good test data to validate the semantics of the models produced [45]. We exemplify these observations by the following case, expanding on Example 1.

Example 2 Consider the relation schema EMPLOYMENT with attribute set Emp, Dept and Mgr and FD set Σ that consists of Emp \rightarrow Dept and Dept \rightarrow Mgr. As in Example 1 assume that only Emp and Mgr have been declared NOT NULL. In this case, the relation

Emp	Dept	Mgr
Hilbert	Math	Gauss
Py thag or as	Math	Gauss
Einstein	Physics	Gauss
Turing	ni	von Neumann
Turing	ni	$G\"odel$

is an Armstrong relation for Σ and the NOT NULL constraints. Note that, according to the semantics under the no information interpretation of the null value **ni** [46, 47, 48], the relation satisfies the FD Dept \rightarrow Mgr since the FD can only be violated if there are distinct tuples that are total and agree in the Dept column, and differ in the Mgr column. When the domain experts inspect this Armstrong relation they are alerted to the fact that Turing has two managers von Neumann and Gödel. More specifically, the domain experts communicate to the team of data engineers that every employee should only occur once in the Emp column. Consequently, the inspection of the Armstrong relation has revealed that under the current design choice neither the meaningful FD Emp \rightarrow Mgr nor the meaningful key Emp have been successfully captured. As a consequence, the team has now a choice of either specifying the FD Emp \rightarrow Mgr additionally to the FDs in Σ , or declaring Dept as NOT NULL.

Example 2 illustrates the assistance that Armstrong relations can provide to data engineers in their task to capture the semantics of the underlying application domain. Unfortunately, the current toolbox of Armstrong relations does not apply to SQL tables where arbitrary attributes can be declared NOT NULL. That means that currently there is no effective support in identifying the set of meaningful FDs over SQL tables. While FDs have been studied in this context [46, 47], the concept of Armstrong relations has not been investigated yet. To address this research gap, we adopt the formal framework of Lien [47], Atzeni and Morfuni [46] who have studied FDs under Zaniolo's no information interpretation of nulls. It is intuitive and well-known that FDs interact quite differently in the presence of NOT NULL constraints than they do over total relations [46]. Therefore, it should come as no surprise that Armstrong relations for FD sets in the presence of NOT NULL constraints can be quite different from Armstrong relations for FDs over total relations. The following example illustrates that currently existing techniques for creating total Armstrong relations do not apply to real SQL table definitions.

Example 3 Consider the relation schema EMPLOYMENT with attributes Emp, Dept and Mgr, and FD set Σ that consists of $Emp \rightarrow Dept$ and $Dept \rightarrow Mgr$. Suppose we would apply the existing techniques to create the total Armstrong relation

Emp	Dept	Mgr
Hilbert	Math	Gauss
Py thag or as	Math	Gauss
Einstein	Physics	Gauss
Turing	Computer Science	von Neumann

for Σ [40]. Then this is only an Armstrong relation for Σ in the special case where all attributes are NOT NULL, but not an Armstrong relation for the SQL table definition

of Example 1. If we incorrectly used this total relation as a concise representation for our SQL table definition of Example 1, then it would give us the false impression that the current design choice successfully captures both the meaningful key Emp and the meaningful FD Emp \rightarrow Mgr. Consequently, there is a real need to study Armstrong relations for FDs in the presence of arbitrary NOT NULL constraints.

Contributions and Organization. In this article we investigate the concept of Armstrong tables for FDs in the presence of NOT NULL constraints. Following previous research [46, 48, 47], we adopt Zaniolo, Lien, Atzeni and Morfuni's *no information* interpretation of null values. This is the most primitive interpretation which allows data engineers to model non-existent information as well as existent information that is currently unknown. In this context, Atzeni and Morfuni have studied FDs in the presence of a null-free subschema (NFS). Essentially, an NFS is the subset of the underlying relation schema whose attributes have been declared NOT NULL. More specifically, Atzeni and Morfuni have presented an axiomatization for the implication of FDs in the presence of an NFS, and an algorithm that decides the corresponding implication problem in time linear in the input size [46]. Note that the combination of FDs and an NFS subsumes the uniqueness and (primary) key constraints over SQL relations.

The objective of this article is to extend the existing theory of Armstrong relations to a toolbox that is applicable to SQL table definitions with functional dependencies. We review previous work in Section 2. Subsequently, we define the underlying concepts and present related previous results in Section 3.

As a first contribution we show in Section 4 that FDs do not enjoy Armstrong relations when null values are permitted to occur. That is, we identify relation schemata and sets of FDs for which no Armstrong relations exist. This is in contrast to the case of total relations. Fortunately, the source of this negative result are so-called non-standard FDs. These are FDs of the form $\emptyset \to A$ with an empty attribute set on the LHS. One may argue that these FDs do not occur in practice since they would enforce all entries in the A-column to be the same. We will show that the class of standard FDs does enjoy Armstrong relations, even in the presence of an arbitrary NFS. The situation is reminiscent of the situation for functional and inclusion dependencies over total relations where only standard FDs and inclusion dependencies enjoy Armstrong databases [49]. For these reasons we will focus on the class of standard FDs in the presence of an NFS.

As a second contribution we characterize Armstrong tables in Section 5. That is, we give sufficient and necessary conditions for a given relation to be an Armstrong table for a given set of standard FDs and a given NFS. Our characterization is based on extensions of the notions of *agree sets* [40], *maximal sets* [42] and "closed" sets from the special case of total relations. However, we demonstrate that in the presence of null values, the closure operation is not idempotent. Therefore, the maximal sets are no longer the intersection generators of closed sets, a property which was fundamental in the development of the results for the special case of total relations [40]. This observation constitutes a significant challenge on the development of our results.

As a third contribution we establish in Section 6 an algorithm that, given a relation schema R, an NFS R_s and a set Σ of standard FDs over R, computes an Armstrong table for Σ and R_s . The algorithm is a non-trivial extension of the algorithm by Mannila and Räihä for FDs over total relations [42]. It is based on the computation of attribute subsets that are maximal for a given attribute, i.e., maximal with the property that the attribute is not functionally dependent on the subset. The computation of the families of maximal attribute subsets is incremental in the given set of standard FDs. The families evolve conceptually differently than they do in the special case of total relations.

As a fourth contribution we investigate the complexity of certain problems associated with Armstrong tables in Section 7. While the problem of finding an Armstrong table remains precisely exponential in general, the size of the Armstrong tables that our algorithm produces is shown to be at most quadratic in the size of a minimum-sized Armstrong table. We show that the size of Armstrong tables for a given set of FDs can be exponentially smaller than an optimal cover of Σ . Therefore, just for the reason of the size of a representation, data engineers should always consider both representations of business rules: as abstract FD sets and as Armstrong tables for these. Finally, we also show that the problem of deciding whether a given attribute set is a Codd key (i.e. it enforces totality and uniqueness) with respect to a given set of FDs and an NFS is NPcomplete. Therefore, our results extend the existing toolbox of Armstrong relations to SQL table definitions with no loss in time efficiency and almost no loss in space efficiency over total relations.

As a fifth contribution we demonstrate in Section 8 how our results carry over to i) schemata on which no Codd key has been specified, and ii) Armstrong tables for standard FDs and an NFS in the world of all FDs. Concerning i) we need to study the combined class of uniqueness constraints and FDs in the presence of an NFS. Concerning ii) Armstrong tables for the world of all FDs must additionally violate all non-trivial non-standard FDs, in contrast to the world of standard FDs.

As a final contribution we illustrate the impact of our results on database practice in Section 9. We show how Armstrong tables can assist data engineers in recognizing meaningful constraints that were previously incorrectly perceived as meaningless. The recognition of these constraints can result in advanced designs that guarantee the absence of data redundancy and update anomalies, ensure new opportunities for the efficient processing of queries, and allow security officers to gain an advanced understanding about possible inference attacks.

In summary, our contributions extend well-known results from total relations to SQL tables. Hence, the resulting toolbox can be applied to instances that occur in real database systems. Moreover, almost all of the nice properties previously established for the special case of total relations [50, 46, 40, 51, 47, 42] are retained in the more general case. We conclude in Section 10 where we also discuss multiple directions of future work.

2 Related Work

Data dependencies have been studied thoroughly in various data models, and we refer the reader to more detailed surveys [2, 52, 53, 54].

2.1 Total relations

We begin with a summary of the work over total relations. Keys and FDs are concepts almost as old as the relational model of data itself [1, 55]. A total relation over a relation schema R satisfies a key X if and only if it satisfies the functional dependency $X \to R$. Therefore, keys are subsumed by FDs. Armstrong established the first axiomatization of FDs over total relations [50], now known as the Armstrong axioms. In fact, Armstrong showed that the Armstrong axioms are even strongly complete for the implication of FDs, i.e., for an arbitrary relation schema and an arbitrary set of FDs on that schema, he constructed a single finite set of tuples which satisfies precisely all implied FDs. That is the reason why such specific relations became known as Armstrong relations. In general, axiomatizations can be applied by data engineers to enumerate all implied data dependencies. In practice, such an enumeration is often desirable, e.g., to validate the correct specification of explicit knowledge, to design and fine-tune databases or to optimize queries. In particular, the completeness of the inference rules ensures that all opportunities of utilizing implicit knowledge for these purposes have been exploited. Furthermore, an analysis of the completeness argument can provide invaluable hints for finding algorithms that efficiently decide the associated implication problem, i.e., to decide for an arbitrarily given FD set $\Sigma \cup \{\varphi\}$ whether Σ implies φ . For FDs over total relations, implication can be decided in time linear in the total number of attributes that occur in the input [56, 57]. Such decision algorithms complement the enumeration algorithm by a further reasoning capability that can make efficient, but only partial decisions about implicit knowledge. These decisions are only partial in the sense that the input to this algorithm must also contain a candidate for an implied functional dependency. In contrast, the enumeration algorithm simply lists all implied dependencies. The reason for the prominence of the implication problem is manifold. An algorithm for testing the implication of dependencies enables us to test whether two given sets of dependencies are equivalent or whether a given set of dependencies is redundant. A solution to these problems is a big step towards automated database schema design [58, 59] which some researchers see as the ultimate goal in dependency theory [60]. Moreover, such an algorithm can be used in relational normalization theory and practice involving many normal form proposals [56, 60, 59, 61, 55], requirements engineering and schema validation [42], data mining [62], in database security [63, 37], view maintenance [17] and in query optimization [34, 35, 64]. More recently, the implication problem has received a lot of attention in other data models as well [65, 66, 46, 67, 68, 69, 70, 71, 72, 73, 74, 75, 76, 77, 78, 79, 80, 81, 15]. New application areas involve data cleaning [82], data transformations [83], consistent query answering [84] and data exchange [85, 86, 87] and data integration [22, 23].

Armstrong relations constitute an invaluable tool for the validation of semantic knowledge, and a user-friendly representation of integrity constraints. Armstrong relations have been deeply studied for keys [51, 88] and FDs [50, 40, 89, 42]. In particular, the existence of Armstrong relations for FDs was shown by Armstrong [50], and Fagin [90] has shown the existence of Armstrong relations for a large class of data dependencies; however, classes of data dependencies do not necessarily enjoy Armstrong relations by any means. The structure of Armstrong relations for the class of FDs over total relations has mainly been investigated by Beeri, Fagin, Statman and Howard [40] and Mannila and Räihä [42]. In the current article, we will extend several of their results from the special case of total to partial relations. The properties of Armstrong relations have also been analyzed for various other classes of data dependencies [26, 91, 43, 49, 92, 42]. An excellent survey on Armstrong databases is [43]. Over the last decades, the concept of Armstrong relations has been found useful in many database applications. For example, informative Armstrong databases are used for the semantic sampling of existing databases [25, 26]. They constitute small subsets of an existing database that satisfy the same data dependencies. In particular, De Marchi and Petit [26] applied the concept of informative Armstrong relations to generate a sample that only used 0.6% of the tuples in an existing real-world database and satisfied the exact same set of data dependencies. The small size of this sample enabled data engineers to identify the FDs meaningful for the application domain. Design aids [41, 42, 93] have been developed that utilize Armstrong databases to help judge, justify, convey, or test the data engineers' understanding of a given relation schema. Recently, empirical evidence has been established that Armstrong relations help data engineers to recognize meaningful FDs that they were not able to recognize without the help of Armstrong relations [44]. Failure to identify some meaningful functional dependency also means that the output of a *requirements analysis* is afflicted with errors. Empirical studies show that more than half the errors which occur during systems development are requirements errors [94, 95]. Requirements errors are also the most common cause of failure in systems development projects [94, 96]. The cost of errors increases exponentially over the development life cycle: it is more than 100 times more costly to correct a defect post-implementation than it is to correct it during requirements analysis [97]. This suggests that it would be more effective to concentrate quality assurance efforts in the requirements analysis stage, in order to catch requirements errors as soon as they occur, or to prevent them from occurring altogether [98]. Hence, there are also strong economic incentives to utilize Armstrong relations for the acquisition of meaningful FDs. Finally, Kolaitis et al. have established first correspondences between unique characterizations of *schema mappings* and the existence of Armstrong bases [39]. An Armstrong basis refers to a finite set of pairs consisting of a source instance and a target instance that is a universal solution for the source instance. Our results may open the way to establish further correspondences.

2.2 Partial relations

We will now comment on some of the work concerning data dependencies in the presence of nulls. One of the most important extensions of Codd's basic relational model [1] is incomplete information. This is mainly due to the high demand for the correct handling of such information in real-world applications. Approaches to deal with incomplete information comprise incomplete relations [99, 100, 101], or-relations [102, 103, 104] or fuzzy relations [77]. Here we focus on incomplete relations.

In the literature, many kinds of null values have been proposed; for example, "missing" or "value unknown at present" [105, 106, 107], "non-existence" [108], "inapplicable" [107], "no information" [109] and "open" [110]. The most primitive is the "no information" interpretation that can be used to model every kind of missing or incomplete information, and its semantics is certainly simple and well understood [46]. Lien [47] investigated FDs in partial relations under this interpretation and established an axiomatization for this class. Interestingly, the transitivity rule, which is part of the Armstrong axioms, is no longer sound in this more general context. Atzeni and Morfuni established axiomatizations and linear-time algorithms for deciding the implication of the combined class of FDs and various existence constraints, including NFSs [46]. It is precisely this line of work we continue in this paper. Recently, Hartmann and Link [48] established an axiomatization and an almost linear-time algorithm for deciding the combined class of functional and multivalued dependencies and NFSs, and showed the equivalence of the implication problem to that of a propositional fragment of Cadoli and Schaerf's S-3 logics [111]. Levene and Loizou introduced and axiomatized the classes of weak and strong FDs with respect to a possible world semantics [76]. The axiomatization of strong FDs is given by the Armstrong axioms, while weak FDs have the same axiomatization as the FDs of Lien [47], Atzeni and Morfuni [46]. However, weak FDs are different from FDs under the no information interpretation. Consider for example the schema EMPLOYMENT with attributes Emp, Dept and Mqr, and the relation r with the two tuples (Turing, Comp. von Neumann) and (Turing, ni, Gödel). The relation satisfies the weak functional dependency $Emp \rightarrow Dept$ since the null value ni can be replaced by Comp (there is a possible world in which $Emp \rightarrow Dept$ is satisfied). However, under the no information interpretation the functional dependency $Emp \rightarrow Dept$ is violated by r. That is, the two tuples have some information on the attribute Emp and the information is the same, but the first tuple has some information for *Dept* while the second tuple has no information for *Dept*. Levene and Loizou also showed that the combined class of weak and strong FDs enjoys Armstrong relations [76]. Paulley axiomatized the implication of a class of strong and lax FDs which advance the query processing in SQL [36]. In particular, Paulley's lax FDs are different from "no information" FDs, e.g. the relation r above satisfies the lax, but not the "no information" FD $Emp \rightarrow Dept$. Another difference is that Paulley accounts for the possible existence of duplicate rows in an SQL table by including tuple identifiers.

Over the last decade, FDs have also received a lot of attention from the XML community [66, 73, 75, 79, 15]. However, the study of Armstrong data trees for XML FDs has only been preliminary so far. In fact, Hartmann et al. [73] characterize the existence of such Armstrong data trees for a particular class of XML FDs under certain assumptions on the underlying XML schema. The results of our present article may provide valuable information for the study of Armstrong data trees for many classes of XML FDs.

3 Preliminaries

In this section we summarize the basic notions required for our treatment of FDs over SQL tables. Our development will follow the extension of Codd's relational model of data [1] to encompass incomplete information by the *no information* null value introduced by Zaniolo [109], applied to FDs by Lien [47], to FDs and null existence constraints by Atzeni and Morfuni [46], and to functional and multivalued dependencies and a null-free subschema by Hartmann and Link [48].

3.1 Total and partial relations

Let $\mathfrak{A} = \{A_1, A_2, \ldots\}$ be a (countably) infinite set of distinct symbols, called attributes. A relation schema is a non-empty, finite subset R of \mathfrak{A} whose attributes represent column names of a relation. Each attribute A of a relation schema R is associated with an infinite domain dom(A) which represents the set of possible values that can occur in the column named A. In order to encompass incomplete information it is assumed that the domain of every attribute has a null value distinct from all the other domain values, denoted by $\mathbf{ni} \in dom(A)$. In the literature, many kinds of null values have been proposed; for example, "missing" or "value unknown at present" [105, 107, 106], "non-existence" [108], "inapplicable" [107], "no information" [109] and "open" [110]. The intention of the null value \mathbf{ni} is to mean "no information". That is, the null value \mathbf{ni} associated with an attribute in a tuple means that no information is available about that attribute for that tuple. This is the most primitive interpretation but can be used to model every kind of missing or incomplete information, and its semantics is certainly simple and well understood [46].

If X and Y are sets of attributes, then we may write XY for $X \cup Y$. If X = $\{A_1,\ldots,A_m\}$, then we may write $A_1\cdots A_m$ for X. In particular, we may write simply A to represent the singleton $\{A\}$. A tuple over R (R-tuple or simply tuple, if R is understood) is a function $t: R \to \bigcup dom(A)$ with $t(A) \in dom(A)$ for all $A \in R$. For $X \subseteq R$ $A \in R$ let t[X] denote the restriction of the tuple t over R to X, and let $dom(X) = \prod_{A \in X} dom(A)$ denote the Cartesian product of the domains of attributes in X. A (partial) relation rover R is a finite set of tuples over R. Let t_1 and t_2 be two tuples over R. It is said that t_1 subsumes t_2 if for every attribute $A \in R$, $t_1[A] = t_2[A]$ or $t_2[A] = ni$ holds. Intuitively, the information content of t_1 subsumes that one of t_2 . In keeping consistency with previous work [46, 47, 112, 109], the following restriction will be imposed: No relation in the database shall contain two tuples t_1 and t_2 such that t_1 subsumes t_2 . Without null values, this amounts to saying that no two tuples are identical, an explicit requirement for database relations. The results we develop in this article for Armstrong tables and the class of FDs, however, remain valid even when we lift this restriction. For a more detailed analysis see Section 8.

For a set $X \subseteq R$, a tuple t over R is X-total, if for all $A \in X$, $t[A] \neq ni$. A relation r over R is X-total, if every tuple t of r is X-total. A relation r over R is a total relation, if it is R-total.

3.2 Agree sets of tuples

Over total relations, Mannila, Räihä, Beeri, Dowd, Fagin and Statman used the notion of an agree set of tuples for the investigation of Armstrong relations. Over partial relations this notion requires a refinement that we introduce in the next definition.

Definition 1 Let R be some relation schema. For two tuples t, t' over R let ag(t, t') = (X, Y) denote the agree set of t and t', where $X = \{A \in R \mid t[A] = t'[A] \land t[A] \neq ni = \neq t'[A]\}$ and $Y = \{A \in R \mid t[A] = t'[A]\}$. For ag(t, t') = (X, Y) we define $ag^s(t, t') = X$ to be the strong agree set, and $ag^w(t, t') = Y$ to be the weak agree set of

t and t'. For a relation r over R let $ag(r) = \{ag(t,t') \mid t,t' \in r \land t \neq t'\}$ denote the agree set of r, let $ag^s(r) = \{X \mid (X,Y) \in ag(r)\}$ denote the strong agree set of r, and let $ag^w(r) = \{Y \mid (X,Y) \in ag(r)\}$ denote the weak agree set of r. Finally, for $X \in ag^s(r)$ let $w(X) = \bigcap\{Y \mid (X,Y) \in ag(r)\}$.

Remark 1 Over total relations [40, 42], it suffices to define $ag(t, t') := \{A \in R \mid t[A] = t'[A]\}$ since a total relation r satisfies $ag^s(r) = ag^w(r)$, and w(X) = X for all $X \in ag^s(r)$. Note that all elements of ag(r) have the form (X, X) in the case of total relations.

Example 4 Let r denote the relation from Example 2:

Emp	Dept	Mgr
Hilbert	Math	Gauss
Py thag or as	Math	Gauss
Einstein	Physics	Gauss
Turing	ni	von Neumann
Turing	ni	$G\ddot{o}del$

The agree set of the last two tuples is

$$({Emp}, {Emp, Dept})$$

and the remaining agree sets of r are:

- $({Dept, Mgr}, {Dept, Mgr})$
- $({Mgr}, {Mgr})$
- *(*Ø,Ø*)*.

Specifically, $w(Emp) = \{Emp, Dept\}.$

3.3 FDs and null-Free subschemata

FDs between sets of attributes have always played a central role in the study of relational databases [56, 59, 1], and seem to be central for the study of database design in other data models as well [66, 71, 75, 76, 78, 10, 79, 80, 81]. In relational databases the notion of a functional dependency is well-understood and the semantic interaction between these dependencies has been syntactically captured by Armstrong's axioms [50]. Lien defined the semantics of a *functional dependency* over partial relations [47]. Here we state this definition in terms of our agree sets.

Definition 2 A functional dependency (FD) over the relation schema R is a statement $X \to Y$ where $X, Y \subseteq R$. The FD $X \to Y$ over R is satisfied by a relation r over R, denoted by $\models_r X \to Y$, if and only if for all $t_1, t_2 \in r$ the following holds: if $X \subseteq ag^s(t_1, t_2)$, then $Y \subseteq ag^w(t_1, t_2)$. FDs of the form $\emptyset \to Y$ are called non-standard, otherwise they are called standard.

Hence, whenever two tuples agree on a non-null restriction to X, then they also agree on the restriction to Y, which may be partial. In other words, an FD $X \to Y$ can only be violated by a relation r if there are two distinct X-total tuples $t_1, t_2 \in r$ such that $t_1[X] = t_2[X]$ and there is some attribute $A \in Y$ such that $t_1[A] \neq t_2[A]$ (i.e. either $t_1[A]$ and $t_2[A]$ have different non-null values, or one of $t_1[A]$ and $t_2[A]$ is null but the other value is non-null).

For total relations the definition of an FD reduces to that of a functional dependency, and so is a correct generalization of the concept. It is also consistent with the no information interpretation [46]. Firstly, tuples with nulls in attributes in X cannot cause a violation of a functional dependency $X \to Y$: the nulls mean that no information is available about those attributes. Secondly, the functional dependence of Y on X forces any two X-total tuples t_1, t_2 where $t_1[X] = t_2[X]$ to have the same information on all the attributes in Y. That is, for all $A \in Y$ we have either $t_1[A] = n\mathbf{i} = t_2[A]$ or $n\mathbf{i} \neq t_1[A] = t_2[A] \neq n\mathbf{i}$. One could also define the semantics differently, e.g. by saying that (i) for all $t_1, t_2 \in r$: if t_1, t_2 are both XY-total and $X \subseteq ag^s(t_1, t_2)$, then $Y \subseteq ag^s(t_1, t_2)$, or (ii) for all $t_1, t_2 \in r$: if $X \subseteq ag^s(t_1, t_2)$, then $Y \subseteq ag^s(t_1, t_2)$. However, option (i) removes the functional dependence of Y on X by stipulating that any tuples that are not Y-total cannot contribute to a violation of the FD $X \to Y$, independently of their values on X. Moreover, option (ii) makes the unnecessarily strict requirement that any tuples t_1, t_2 with $X \subseteq ag^s(t_1, t_2)$ must also be Y-total.

According to Atzeni and Morfuni [46], a null-free subschema (NFS) over the relation schema R is a an expression R_s where $R_s \subseteq R$. The NFS R_s over R is satisfied by a partial relation r over R, denoted by $\models_r R_s$, if and only if for all $t \in r$ and for all $A \in R_s$ we have $t[A] \neq ni$. Hence, the satisfaction of a null-free subschema R_s over R requires that partial relations over R are R_s -total. Without loss of generality, we can assume that on each relation schema a single NFS is defined. Null-free subschemata occur naturally in database practice: SQL allows one to specify attributes as NOT NULL, cf. Example 1. Consequently, the set of such attributes over a table definition would form the single null-free subschema over this table.

For a set Σ of constraints over some relation schema R, we say that a (partial) relation r over R satisfies Σ , denoted by $\models_r \Sigma$, if r satisfies every element of Σ . If for some $\sigma \in \Sigma$ the relation r does not satisfy σ we sometimes say that r violates σ (in which case r also violates Σ) and write $\not\models_r \sigma$ ($\not\models_r \Sigma$).

Example 5 For Example 1 we obtain the relation schema

 $EMPLOYMENT = \{Emp, Dept, Mgr\}$

with NFS EMPLOYMENT_s = {Emp, Mgr}. The relation

Emp	Dept	Mgr
Turing	ni	von Neumann
Turing	ni	$G\ddot{o}del$

satisfies EMPLOYMENT_s as well as the FDs $Emp \to Dept$ and $Dept \to Mgr$. However, the relation violates the FD $Emp \to Mgr$.

3.4 Implication and inference

For the design, maintenance and applications of a relational database, data dependencies are identified as semantic constraints on the relations which are intended to be instances of the database schema. During the design process or lifetime of a database one usually needs to determine further dependencies which are logically implied by the given ones. In line with the literature of database constraints, we restrict our attention to the implication of constraints in some fixed class C: FDs in the presence of an NFS.

Let R be a relation schema, let $R_s \subseteq R$ denote an NFS over R, and let $\Sigma \cup \{\varphi\}$ be a set of FDs over R. We say that Σ *implies* φ in the presence of R_s , denoted by $\Sigma \models_{R_s} \varphi$, if every relation r over R that satisfies Σ and R_s also satisfies φ . If Σ does not imply φ in the presence of R_s we may also write $\Sigma \not\models_{R_s} \varphi$.

The implication problem for functional dependencies in the presence of a null-free subschema is to decide, given any relation schema R, any NFS R_s over R, and any set $\Sigma \cup \{\varphi\}$ of FDs over R, whether $\Sigma \models_{R_s} \varphi$. For the class of FDs in the presence of an NFS, the sets $\Sigma \cup \{\varphi\}$ over a relation schema R are always finite. Moreover, if $R_s = \emptyset$ we also write $\Sigma \models \varphi$ instead of $\Sigma \models_{\emptyset} \varphi$. This covers the case where every attribute is NULL [47]. The case where every attribute is NOT NULL is covered when $R_s = R$.

Note that for FDs (in the presence of an NFS) it does not matter whether we restrict our attention to relations that are finite, i.e., the implication problem coincides with the finite implication problem where only finite relations are considered [2]. For this reason, we will only speak of the implication problem.

For an FD set Σ over a relation schema R and an NFS R_s over R, let the FD set $\Sigma_{R_s}^* = \{\varphi \mid \Sigma \models_{R_s} \varphi\}$ denote the *semantic closure* of Σ . For a finite FD set $\Sigma \cup \{\varphi\}$ and a set \mathfrak{R} of inference rules let $\Sigma \vdash_{\mathfrak{R}} \varphi$ denote an *inference* of φ from Σ by \mathfrak{R} . That is, there is some sequence $\gamma = [\sigma_1, \ldots, \sigma_n]$ of FDs such that $\sigma_n = \varphi$ and every σ_i is an element of Σ or results from an application of an inference rule in \mathfrak{R} to some FDs in $\{\sigma_1, \ldots, \sigma_{i-1}\}$. For a finite FD set Σ let $\Sigma_{\mathfrak{R}}^+ = \{\varphi \mid \Sigma \vdash_{\mathfrak{R}} \varphi\}$ denote the *syntactic closure* of Σ under inferences by \mathfrak{R} . \mathfrak{R} is said to be *sound* (*complete*) for the implication of FDs in the presence of an NFS if for every relation schema R, for every NFS R_s over R and for every FD set Σ over R we have $\Sigma_{\mathfrak{R}}^+ \subseteq \Sigma_{R_s}^* (\Sigma_{R_s}^* \subseteq \Sigma_{\mathfrak{R}}^+)$. The (finite) set \mathfrak{R} is said to be a (finite) axiomatization for the implication of FDs in the presence of an NFS if \mathfrak{R} is both sound and complete for the implication of FDs in the presence of an NFS.

Example 6 Consider the relation schema EMPLOYMENT with the NFS EMPLOYMENT_s from Example 5 again. Let Σ denote the set of FDs over EMPLOYMENT that consists of Emp \rightarrow Dept and Dept \rightarrow Mgr. Since the relation in Example 5 satisfies the NFS EMPLOYMENT_s as well as Σ , but it violates the FD Emp \rightarrow Mgr we conclude that Emp \rightarrow Mgr is not in the semantic closure $\Sigma^*_{\text{EMPLOYMENT_s}}$ of Σ .

Remark 2 It was shown recently [48] that an FD set Σ implies an FD φ in the presence of the NFS R_s if and only if the set Σ' of Horn clauses resulting from Σ logically implies the Horn clause φ' resulting from φ in Cadoli and Schaerf's para-consistent family of S-3 logics. Here, S consists of those propositional variables that correspond to the attributes of the NFS R_s and must be interpreted as either true or false, whereas variables outside of S can also be interpreted as both true and false. Importantly, the truth tables that

Table 1: Axiomatization of FDs in the presence of the NFS R_s .

define the logical operators of negation, disjunction and conjunction are the tables that Codd suggested for extending the relational algebra by means of a three-valued logic and the null substitution principal, and which has been adopted by SQL [99].

Atzeni and Morfuni have established a finite axiomatization [46] for the implication of FDs in the presence of an NFS. The inference rules are given in Table 1.

Note that the null transitivity rule can only infer the FD $X \to Z$ from the FDs $X \to Y$ and $Y \to Z$, if all the attributes in Y - X have been declared NOT NULL, i.e., are members of the NFS R_s . Also note that the so-called *augmentation rule*

$$\frac{X \to Y}{XZ \to Y}$$

follows from the reflexivity axiom and the null transitivity rule [46].

Example 7 Consider the relation schema EMPLOYMENT with the NFS EMPLOYMENT_s from Example 5 again. Then the FD $Emp \rightarrow Mgr$ cannot be inferred from the two FDs $Emp \rightarrow Dept$ and $Dept \rightarrow Mgr$ by means of the null transitivity rule since the attribute Dept is not an element of EMPLOYMENT_s.

Atzeni and Morfuni also established a linear-time algorithm for deciding the implication problem for FDs in the presence of an NFS[46]. As Beeri and Bernstein did for total relations [56], Atzeni and Morfuni utilized the notion of an *attribute closure*

$$X^*_{\Sigma,R_s} = \{ A \in R \mid \Sigma \models_{R_s} X \to A \}$$

of an attribute set X with respect to an FD set Σ and an NFS R_s over the relation schema R [46]. An FD X \rightarrow Y over R is implied by Σ in the presence of R_s if and only if $Y \subseteq X^*_{\Sigma,R_s}$ holds [46]. Algorithm 1 on page 15 computes the attribute closure X^*_{Σ,R_s} of X with respect to Σ and R_s over R [46].

Algorithm 1 corrects the algorithm originally proposed in [46]. For example, consider the relation schema EMPLOYMENT with the NFS $\{Dept\}$ and let Σ consist of the two FDs $Emp \rightarrow Dept$ and $Dept \rightarrow Mgr$. On input

$$({Emp}, \Sigma, {Dept}, EMPLOYMENT)$$

Algorithm 1 Closure Computation

```
1: procedure NFSCLOSURE(X, \Sigma, R_s, R)
 2:
        Closure \leftarrow X;
 3:
        repeat
            OldClosure \leftarrow Closure;
 4:
            for all V \to W \in \Sigma do
 5:
                if V \subseteq Closure \cap XR_s then
 6:
                    Closure \leftarrow Closure \cup W;
 7:
                end if
 8:
            end for
 9:
        until Closure=OldClosure;
10:
        return(Closure);
11:
12: end procedure
```

the original algorithm [46] returns the set $\{Emp, Dept\}$, but the correct result is the set $\{Emp, Dept, Mgr\}$.

Example 8 Consider the relation schema EMPLOYMENT with the NFS EMPLOYMENT_s from Example 5 again, and let Σ consist of the two FDs Emp \rightarrow Dept and Dept \rightarrow Mgr. On input

```
(\{Emp\}, \Sigma, EMPLOYMENT_s, EMPLOYMENT)
```

Algorithm 1 returns the set $\{Emp, Dept\}$. Consequently, the FD $Emp \rightarrow Mgr$ is not implied by Σ in the presence of the NFS EMPLOYMENT_s since Mgr is not an element of $\{Emp\}_{\Sigma, EMPLOYMENT_s}^* = \{Emp, Dept\}$.

3.5 Uniqueness constraints and Codd keys

Functional dependencies in the presence of arbitrary null-free subschemata subsume SQL's uniqueness constraint as well as so-called Codd keys.

A uniqueness constraint over a relation schema R is an expression unique(X) with $X \subseteq R$. A relation r over R satisfies unique(X) if and only if for all distinct $t, t' \in r$ the following holds: if t and t' are X-total, then $t[X] \neq t'[X]$. It follows that r satisfies unique(X) if and only if r satisfies the FD $X \to R$.

Moreover, a *Codd key* over a relation schema R is an expression Codd(X) with $X \subseteq R$. A relation r over R satisfies Codd(X) if and only if r is X-total and for all distinct $t, t' \in r$ it is true that $t[X] \neq t'[X]$. From this definition it follows that r satisfies Codd(X) if and only if r satisfies the NFS X and satisfies the FD $X \to R$.

It is now relatively easy to see that we can express uniqueness constraints and Codd keys by means of FDs and an NFS.

Theorem 1 Let Σ be a set of FDs and let R_s be an NFS over R. Then the following hold:

1. $\Sigma \models_{R_s} unique(X)$ if and only if $\Sigma \models_{R_s} X \to R$,

2. $\Sigma \models_{R_s} Codd(X)$ if and only if $\Sigma \models_{R_s} X \to R$ and $X \subseteq R_s$.

Proof The first equivalence is a straight consequence of the fact that a relation satisfies unique(X) if and only if it satisfies the FD $X \to R$. For the second equivalence we first show that if $\Sigma \models_{R_s} X \to R$ and $X \subseteq R_s$ hold, then $\Sigma \models_{R_s} Codd(X)$ holds as well. In fact, let r denote an arbitrary relation over R that satisfies Σ and R_s . It follows that r satisfies $X \to R$ as well since $\Sigma \models_{R_s} X \to R$. Since $X \subseteq R_s$ holds, we know that r satisfies the NFS X. Consequently, r satisfies Codd(X).

It remains to show that if $\Sigma \not\models_{R_s} X \to R$ or $X \not\subseteq R_s$, then $\Sigma \not\models_{R_s} Codd(X)$.

Suppose first that $\Sigma \not\models_{R_s} X \to R$. Then there is some relation r over R that satisfies Σ and the NFS R_s , but violates the FD $X \to R$. Consequently, there are two tuples $t, t' \in r$ such that t[X] = t'[X], t and t' are X-total and $t \neq t'$. The two-tuple relation $r' = \{t, t'\}$ shows that $\Sigma \not\models_{R_s} Codd(X)$ since r' satisfies Σ and R_s (since $r' \subseteq r$ holds), but violates Codd(X).

Suppose now that $X \not\subseteq R_s$ holds, i.e., $X - R_s$ is non-empty. In this case we define a single-tuple relation $r := \{t\}$ for some tuple t over R such that t[A] := ni for all $A \in R - R_s$ and $t[B] \in dom(B) - \{ni\}$ for all $B \in R_s$. It follows that r satisfies Σ and the NFS R_s , but r violates Codd(X) since $X \cap (R - R_s) \neq \emptyset$.

3.6 Armstrong relations

Let Σ be a set of constraints in class C over some relation schema R. A relation r over R is said to be an Armstrong relation for Σ , if for all φ in C over R it is true that r satisfies φ if and only if Σ implies φ . Hence, r is a perfect representation of the constraint set Σ in the sense that it satisfies all the constraints implied by Σ , but violates all the constraints not implied by Σ .

A class \mathcal{C} of constraints is said to *enjoy Armstrong relations* if and only if for every relation schema R, and for every set Σ of constraints in \mathcal{C} over R there is some relation r over R that is an Armstrong relation for Σ .

In this paper, we are interested in the class C of FDs in the presence of an NFS, which permits arbitrary FD sets together with a single arbitrary null-free subschema. In this context, we also speak of *Armstrong tables* instead of Armstrong relations. Note that an NFS R_s implies another NFS R'_s if and only if $R'_s \subseteq R_s$ holds.

Example 9 The relation of Example 3 is an Armstrong table for the FD set

$$\Sigma = \{ Emp \to Dept, Dept \to Mgr \}$$

and the NFS $R_s = R = \{Emp, Dept, Mgr\}$. The relation of Example 2 is an Armstrong table for the FD set Σ and the NFS $R_s = \{Emp, Mgr\}$ over $R = \{Emp, Dept, Mgr\}$. The relation of Example 2 is not an Armstrong table for Σ and $R_s = R$, and the relation of Example 3 is not an Armstrong table for Σ and $R_s = \{Emp, Mgr\}$: it satisfies $R'_s = R$ and the FD Emp \rightarrow Mgr, but neither is implied by Σ and R_s .

Further Outline. The development of these notions have led us to several questions. A first fundamental question is whether the class of FDs in the presence of an NFS enjoys Armstrong tables. We will pursue a first answer to this question in Section 4. Another question is the following: how can we validate that the relations in Examples 2 and 3 are indeed Armstrong tables? More generally, we would like to have a simple characterization that allows us to judge whether any given relation is an Armstrong table for an arbitrarily given set of FDs and an arbitrarily given NFS R_s . Such characterizations will be established in Section 5. These will assist us in Section 6 when we develop an algorithm that computes an Armstrong table for an arbitrarily given NFS. We investigate questions regarding the complexity of Armstrong tables in Section 7.

4 FDs Do Not Enjoy Armstrong Tables

For total relations, i.e. in the special case where the null-free subschema contains all the attributes of the underlying relation schema, it is well-known that FDs enjoy Armstrong relations [50, 40, 43]. We will show now, however, that the class of FDs does not enjoy Armstrong tables if there are attributes that are not declared NOT NULL. Intuitively in such a case, non-standard FDs force all the tuples of a relation to have the same value on some attribute. For relations to be Armstrong, however, it may also be required that the values of such an attribute do differ for some distinct tuples. Consequently, Armstrong tables cannot exist in general.

Theorem 2 The class of functional dependencies in the presence of a null-free subschema does not enjoy Armstrong tables.

Proof We show that there is a relation schema R, an NFS R_s and an FD set Σ over R such that there is no Armstrong table for Σ and R_s .

Let R = ABC, $R_s = BC$ and let $\Sigma = \{\emptyset \to A, A \to B\}$. First we show that an Armstrong table r for Σ and R_s must violate the FDs $AB \to C$ and $C \to B$. In fact, the relation

A	B	C
a	b	С
a	b	c'

satisfies the FD set Σ and the NFS R_s , but violates the FD $AB \to C$. Consequently, $\Sigma \not\models_{R_s} AB \to C$. Moreover, the relation

$$\begin{array}{c|c|c} A & B & C \\ \hline \texttt{ni} & b' & c \\ \texttt{ni} & b & c \\ \end{array}$$

satisfies the FD set Σ and the NFS R_s , but violates the FD $C \to B$. Consequently, $\Sigma \not\models_{R_s} C \to B$. We conclude that an Armstrong table for Σ in the presence of R_s must at least satisfy Σ and violate the FDs $AB \to C$ and $C \to B$.

We show now that any relation that satisfies the FD $\emptyset \to A$ and violates the FDs $AB \to C$ and $C \to B$ must violate the FD $A \to B$. Let r be a relation that satisfies the

FD $\emptyset \to A$ and violates the FDs $AB \to C$ and $C \to B$. Since r violates $AB \to C$ there are two distinct tuples $t_1, t_2 \in r$ such that $t_1[AB] = t_2[AB]$ and t_1, t_2 are both AB-total, and $t_1[C] \neq t_2[C]$. Since r violates the FD $C \to B$ there are two distinct tuples $t_3, t_4 \in r$ such that $t_3[C] = t_4[C]$ and t_3, t_4 are both C-total, and $t_3[B] \neq t_4[B]$. Since r satisfies the FD $\emptyset \to A$, it follows that $t_3[A] = t_4[A] = t_1[A]$, and, in particular, that t_3 and t_4 are A-total. Hence, the tuples t_3, t_4 witness that r violates the FD $A \to B$. Consequently, r cannot be an Armstrong table for Σ and R_s .

For the proof of Theorem 2 it actually suffices to consider the FD $AB \to C$ and the null-free subschema $R_s = BC$. A relation that violates the FD $AB \to C$ must contain two tuples that are A-total. A relation that is Armstrong for Σ and R_s should contain a tuple that carries the no information null value ni on A. Consequently, no relation can violate the FD $AB \to C$ and the NFS A, and satisfy the FD $\emptyset \to A$. Hence, there is no Armstrong table for Σ and R_s .

However, the proof of Theorem 2 also shows that Lien's class of FDs over partial relations [47] does not enjoy Armstrong relations. That is, the proof argument also applies when no NFS R_s is present. In particular, note that $\Sigma \not\models AB \rightarrow C$ and $\Sigma \not\models C \rightarrow B$ also hold.

Corollary 1 Lien's class of FDs does not enjoy Armstrong relations.

The next two examples illustrate the difference between the special case of total relations (Example 10) and the general case (Example 11) we consider here.

Example 10 Consider the relation schema

EMPLOYMENT =
$$\{Emp, Dept, Mgr\},\$$

the NFS $R_s = \{Emp, Dept, Mgr\}$ and let Σ consist of the two FDs $\emptyset \to Mgr$ and $Mgr \to Dept$. The relation r

Emp	Dept	Mgr
Hilbert	Maths	Gauss
Py thag or as	Maths	Gauss

is an Armstrong table for Σ and R_s . In particular, note that the FD $\emptyset \to$ Dept is implied by Σ and R_s .

Example 11 Consider the relation schema

 $EMPLOYMENT = \{Emp, Dept, Mgr\},\$

the NFS $R_s = \{Emp, Dept\}$ and let Σ consist again of the two FDs $\emptyset \to Mgr$ and $Mgr \to Dept$. This is the case of the proof in Theorem 2. Here, the relation r'

	Emp	Dept	Mgr
-	Gauss	Maths	ni
	Gauss	Physics	ni
	Turing	Computers	$G\ddot{o}del$
	von Neumann	Computers	Gödel

is not an Armstrong table for Σ and R_s . Indeed, it violates all FDs not implied by Σ and R_s , and satisfies all FDs implied by Σ and R_s except $\emptyset \to Mgr$.

While the result of Theorem 2 is negative in general, we will see that Armstrong tables do exist for the class of standard FDs in the presence of an NFS. Since it can be argued that non-standard FDs rarely occur in practice, the situation is actually rather pleasant. It is reminiscent of Fagin and Vardi's result that the class of standard FDs and inclusion dependencies over total relations enjoys Armstrong databases, while the class of FDs and inclusion dependencies does not [49].

5 Characterization of Armstrong Tables

In this section we continue Lien [47], Atzeni and Morfuni's study [46] of the class of FDs in the context of partial relations and NFSs, respectively. As a first main result we extend Mannila, Räihä, Beeri, Dowd, Fagin and Statman's characterization of Armstrong relations for FDs from total relations [40, 42]. Our generalization requires an extension of Demetrovics' notion of maximal sets of attributes [51, 42] and uses a refinement of Mannila, Räihä, Beeri, Dowd, Fagin and Statman's notion of an agree set of tuples, already given in Definition 1. Due to Theorem 2 we will focus in Sections 5, 6 and 7 on the class \mathcal{C} of standard FDs. In Section 8 we will look at extensions of our results.

5.1 Maximal families of sets

Definition 3 Let Σ be a set of standard FDs and let R_s be an NFS over a relation schema R. For an attribute $A \in R$ we define the maximal set $max_{\Sigma,R_s}(A)$ of A with respect to Σ and R_s as follows:

$$\max_{\Sigma,R_s}(A) := \{ \emptyset \neq X \subseteq R \mid \Sigma \not\models_{R_s} X \to A \land \\ \forall B \in R - X(\Sigma \models_{R_s} XB \to A) \}.$$

The maximal sets of R with respect to Σ and R_s are $max_{\Sigma,R_s}(R) = \bigcup_{A \in R} max_{\Sigma,R_s}(A)$. If Σ and R_s are clear from the context, we may simply write max(A) and max(R), respectively.

Thus, the maximal sets of an attribute $A \in R$ with respect to Σ and R_s are the maximal attribute subsets of R that do not functionally determine A.

Example 12 Let $\Sigma = \{Emp \rightarrow Dept, Dept \rightarrow Mgr\}$ and let $R_s = \{Emp, Dept, Mgr\}$ be an NFS over the relation schema EMPLOYMENT. Then $max_{\Sigma,R_s}(Emp)$ consists of $\{Dept, Mgr\}, max_{\Sigma,R_s}(Dept)$ contains the single element $\{Mgr\}, and max_{\Sigma,R_s}(Mgr) = \emptyset$.

For $R_s = \{Emp, Mgr\}, max_{\Sigma,R_s}(Emp) \text{ consists of } \{Dept, Mgr\}, max_{\Sigma,R_s}(Dept) \text{ consists of } \{Mgr\}, and max_{\Sigma,R_s}(Mgr) \text{ consists of } \{Emp\}.$

5.2 Characterizations

For a relation r over a relation schema R let

$$total(r) := \{ A \in R \mid \forall t \in r(t[A] \neq \mathtt{ni}) \}$$

denote the set of those attributes A of R such that no tuple $t \in r$ carries a null value ni on A.

Theorem 3 Let R be some relation schema, let Σ be a set of standard FDs and let R_s be an NFS over R. For all relations r over R it holds that r is an Armstrong table for Σ and R_s if and only if both of the following conditions are satisfied:

1. for all non-empty $X \subseteq R$ we have

$$X^*_{\Sigma,R_s} = \bigcap \{ w(Z) \mid X \subseteq Z \in ag^s(r) \}, and$$

2. $total(r) = R_s$.

Proof Sufficiency. Let r be a relation over R that satisfies the conditions. We show that r is an Armstrong table for Σ and R_s .

Let $X \to A \in \Sigma$. That is, $A \in X^*_{\Sigma,R_s}$. Assume that there are distinct tuples $t, t' \in r$ such that t[X] = t'[X] and t, t' are X-total. That is, $X \subseteq X' = ag^s(t, t')$. Hence, the first condition shows that $A \in w(X')$, and thus $A \in ag^w(t, t')$. Therefore, t[A] = t'[A]holds. We have shown that r satisfies Σ .

Let $X \to A \notin \Sigma_{R_s}^*$. Hence, $A \notin X_{\Sigma,R_s}^*$. By the first condition there is some $Z \in ag^s(r)$ such that $X \subseteq Z$ and $A \notin w(Z)$. In particular, there are tuples t, t' such that $X \subseteq Z = ag^s(t, t')$ and $A \notin ag^w(t, t')$. That is, we have t[X] = t'[X], t, t' are X-total and $t[A] \neq t'[A]$. This shows that r violates every FD not in $\Sigma_{R_s}^*$.

Finally, the condition $total(r) = R_s$ ensures that r satisfies every NFS implied by R_s and violates every NFS not implied by R_s . Consequently, r is indeed an Armstrong table for Σ and R_s .

Necessity. Let r be a relation over R that is Armstrong for Σ and R_s . We show that r satisfies the conditions. Let $t, t' \in r$ be distinct tuples such that $X \subseteq X' = ag^s(t, t')$. As r satisfies $\Sigma_{R_s}^*$ we have $X_{\Sigma,R_s}^* \subseteq ag^w(t,t')$, and thus $X_{\Sigma,R_s}^* \subseteq w(X')$. Therefore, $X_{\Sigma,R_s}^* \subseteq \bigcap\{w(Z) \mid X \subseteq Z \in ag^s(r)\}$ holds.

Next we show that $X_{\Sigma,R_s}^* \supseteq \bigcap \{w(Z) \mid X \subseteq Z \in ag^s(r)\}$. Assume there is an $A \notin X_{\Sigma,R_s}^*$ such that $A \in \{w(Z) \mid X \subseteq Z \in ag^s(r)\}$. Then we have $A \in ag^w(t,t')$ for all distinct tuples $t, t' \in r$ with $X \subseteq Z = ag^s(t,t')$. That is, r satisfies $X \to A$. This, however, contradicts the assumption $A \notin X_{\Sigma,R_s}^*$ since r is Armstrong for Σ and R_s . Consequently, $X_{\Sigma,R_s}^* \supseteq \bigcap \{w(Z) \mid X \subseteq Z \in ag^s(r)\}$ holds. Finally, since r is Armstrong for Σ and R_s .

For the special case of total relations, i.e. where $R_s = R$, it is well-known that every FD set Σ and the NFS R over R defines a closure operator $(\cdot)_{\Sigma,R}^* : \mathcal{P}(R) \to \mathcal{P}(R)$ by mapping every attribute subset $X \subseteq R$ to its attribute closure $X_{\Sigma,R}^*$. In fact, $(\cdot)_{\Sigma,R_s}^*$ is a closure operator since it is extensive $(X \subseteq X_{\Sigma,R}^*)$, increasing (for $X \subseteq Y$ we have $X_{\Sigma,R}^* \subseteq Y_{\Sigma,R}^*$) and idempotent $\left(\left(X_{\Sigma,R}^* \right)_{\Sigma,R}^* = X_{\Sigma,R}^* \right)$. In particular, an attribute subset X of R is said to be closed with respect to the FD set Σ and the NFS R if $X_{\Sigma,R}^* = X$. The set of all subsets of R that is closed with respect to the FD set Σ and the NFS R if $X_{\Sigma,R}^* = X$. The set of all subsets of R that is closed with respect to the FD set Σ and the NFS R is denoted by $cl_{\Sigma,R}(R)$. For the special case where $R_s = R$ it is well-known that $cl_{\Sigma,R_s}(R)$ is closed under intersection. Thus there is a unique minimal subfamily of generators $gen_{\Sigma,R}(R) \subseteq cl_{\Sigma,R}(R)$ such that each member of $cl_{\Sigma,R}(R)$ can be expressed as an intersection of sets in $gen_{\Sigma,R}(R)$ [40].

Beeri, Dowd, Fagin, and Statman [40] have shown that r is an Armstrong table for an FD set Σ and the NFS R over R if and only if $gen_{\Sigma,R}(R) \subseteq ag^s(r) \subseteq cl_{\Sigma,R}(R)$ holds. Later, Mannila and Räihä [42] have shown that for an arbitrary relation schema R, an arbitrary FD set Σ and the special NFS R over R it is true that $max_{\Sigma,R}(R) = gen_{\Sigma,R}(R)$.

If we allow arbitrary null-free subschemata R_s of R, then the situation is different. In particular, $(\cdot)_{\Sigma,R_s}^* : \mathcal{P}(R) \to \mathcal{P}(R)$ is no longer idempotent, and therefore no closure operator, as the next example illustrates.

Example 13 Let R denote the relation schema EMPLOYMENT from Example 2 where the FD set Σ consists of the two FDs Emp \rightarrow Dept and Dept \rightarrow Mgr, and the NFS R_s consists of the attributes Emp and Mgr. Then the attribute subset closures are:

- $\emptyset_{\Sigma,R_s}^* = \emptyset$,
- $\{Emp\}_{\Sigma,R_{s}}^{*} = \{Emp, Dept\},\$
- $\{Dept\}_{\Sigma,R_s}^* = \{Dept, Mgr\},\$
- $\{Mgr\}_{\Sigma,R_s}^* = \{Mgr\},\$
- $\{Emp, Dept\}_{\Sigma, R_s}^* = R,$
- $\{Emp, Mgr\}_{\Sigma, R_s}^* = R,$
- ${Dept, Mgr}_{\Sigma,R_s}^* = {Dept, Mgr}, and$
- $\{Emp, Dept, Mgr\}_{\Sigma, R_s}^* = R.$

Hence, we can see that the attribute subset closure $(\cdot)_{\Sigma,R_s}^*$ is not idempotent, and

$$cl_{\Sigma,R_s}(R) = \{\emptyset, \{Mgr\}, \{Dept, Mgr\}\} = gen_{\Sigma,R_s}(R),$$

is different from

$$max_{\Sigma,R_s}(R) = \{ \{ Emp \}, \{ Mgr \}, \{ Dept, Mgr \} \}.$$

Since cl_{Σ,R_s} does not define a closure operator for arbitrary NFS R_s over R the concept of *closed attribute subsets* is no longer useful in this context. For this reason, we will utilize the maximal set families $max_{\Sigma,R_s}(R)$ to characterize the situation when an arbitrary relation is an Armstrong table for a given FD set Σ and a given NFS R_s over the relation schema R. **Theorem 4** Let R be some relation schema, let Σ be a set of standard FDs and let R_s be an NFS over R. For all relations r over R it holds that r is an Armstrong table for Σ and R_s if and only if the following conditions are satisfied:

- 1. $\forall A \in R \forall X \in max_{\Sigma,R_s}(A) (X \in ag^s(r) \land A \notin w(X)),$
- 2. $\forall X \in ag^s(r)(X^*_{\Sigma,R_s} \subseteq w(X)), and$
- 3. $total(r) = R_s$.

Proof Sufficiency. Let r be some relation over R that satisfies conditions 1., 2. and 3. We show that r is an Armstrong table for Σ and R_s .

Let $X \to A \in \Sigma$. Assume that there are distinct $t, t' \in r$ such that t[X] = t'[X]and t is X-total. That is, $X \subseteq X' = ag^s(t, t')$. Note that $A \in (X')_{\Sigma,R_s}^*$ by soundness of the augmentation rule. Hence, condition 2. implies that $A \in w(X')$. In particular, $A \in ag^w(t, t')$. Therefore, t[A] = t'[A]. Hence, r satisfies Σ .

Let $X \to A \notin \Sigma_{R_s}^*$. It follows that there is some $X' \in max_{\Sigma,R_s}(A)$ such that $X \subseteq X'$ and $A \notin (X')_{\Sigma,R_s}^*$. Condition 1. implies that $X' \in ag^s(r)$ and $A \notin w(X')$. Hence, there is some $Y \in ag^w(r)$ such that $(X',Y) \in ag(r)$ and $A \notin Y$. This shows that there are two distinct $t, t' \in r$ such that t[X'] = t'[X'] and t, t' are X'-total and $t[A] \neq t'[A]$. We have shown that r violates every functional dependency that is not in $\Sigma_{R_s}^*$.

Condition 3. ensures that r satisfies every NFS implied by R_s , and violates every NFS not implied by R_s . Consequently, r is an Armstrong table for Σ and R_s .

Necessity. Let r be some relation over R that is an Armstrong table for Σ and R_s . We show that r satisfies conditions 1., 2. and 3.

Let $A \in R$, and let $X \in max_{\Sigma,R_s}(A)$. That is, $\Sigma \not\models_{R_s} X \to A$ and for all $B \in R - X$ it is true that $\Sigma \models_{R_s} XB \to A$. Since r is an Armstrong table for Σ and R_s it follows that r violates $X \to A$ and for all $B \in R - X$ that r satisfies the FD $XB \to A$. The violation of $X \to A$ implies that there are distinct $t, t' \in r$ such that $X \subseteq ag^s(t, t')$ and $A \notin ag^w(t, t')$. If there was some attribute C of R in $ag^s(t, t') - X$, then r would violate the FD $XC \to A$. Consequently, $X = ag^s(t, t')$. We have just shown that for every $A \in R$ and for every $X \in max_{\Sigma,R_s}(A)$ it is true that $X \in ag^s(r)$ and $A \notin w(X)$, i.e., condition 1. holds.

Next we show that r satisfies condition 2. Therefore, let $X \in ag^s(r)$. We need to show that $X_{\Sigma,R_s}^* \subseteq w(X)$. Let A be some attribute of R such that $A \notin w(X)$. That is, there is some $Y \in ag^w(r)$ such that $(X,Y) \in ag(r)$ and $A \notin Y$. Consequently, there are some distinct $t, t' \in r$ such that $X = ag^s(t, t')$ and $A \notin ag^w(t, t')$. That is, r violates the FD $X \to A$. Since r is an Armstrong table for Σ and R_s it follows that $A \notin X_{\Sigma,R_s}^*$. We have just shown that $X_{\Sigma,R_s}^* \subseteq w(X)$.

Since r is an Armstrong table for Σ and R_s it follows that $total(r) = R_s$.

Remark 3 Theorem 4 generalizes the characterization of Armstrong relations for standard FDs from the special case where the NFS R_s is R. In fact, if r is total, then condition 1. says that for all $A \in R$ and for all $X \in \max_{\Sigma,R_s}(A)$ we have $X \in ag^s(r)$ (and that $A \notin X$ which follows from $X \in \max_{\Sigma,R_s}(A)$). That is, condition 1. says that $\max_{\Sigma,R_s}(R) \subseteq ag^s(r)$. Furthermore, condition 2. says that $\forall X \in ag^s(r)(X_{\Sigma}^* \subseteq X)$, i.e., all (strong) agree sets are closed: $ag^s(r) \subseteq cl_{\Sigma,R_s}(R)$. The following example illustrates Theorem 4.

Example 14 Let EMPLOYMENT and the FD set

$$\Sigma = \{ Emp \to Dept, Dept \to Mgr \}$$

as before, but let $R_s = \{Emp, Mgr\}$. The following relation r:

Emp	Dept	Mgr
Hilbert	Math	Gauss
Py thag or as	Math	Gauss
Einstein	Physics	Gauss
Turing	ni	von Neumann
Turing	ni	$G\ddot{o}del$

is an Armstrong table for Σ and R_s . Indeed, the maximal set families are:

- $max_{\Sigma,R_s}(Emp) = \{\{Dept, Mgr\}\},\$
- $max_{\Sigma,R_s}(Dept) = \{\{Mgr\}\}, and$
- $max_{\Sigma,R_s}(Mgr) = \{ \{ Emp \} \}.$

The agree sets of r are:

- (Ø,Ø),
- $({Dept, Mgr}, {Dept, Mgr}),$
- $({Mgr}, {Mgr}),$
- $({Emp}, {Emp, Dept})$.

The conditions 1., 2. and 3. of Theorem 4 are satisfied.

One may expect condition 2. of Theorem 4 to read

$$\forall X \in ag^s(r)(X^*_{\Sigma,R_s} = w(X)),$$

but a slight change of Example 14 shows that this condition would fail to recognize some Armstrong tables.

Example 15 Let EMPLOYMENT and

 $\Sigma = \{ Emp \to Dept, Dept \to Mgr \}$

be as before, but let $R_s = \{Mgr\}$. The following relation r:

Emp	Dept	Mgr
Hilbert	Math	Gauss
Py thag or as	Math	Gauss
ni	Astronomy	Newton
ni	Physics	Newton
Turing	ni	von Neumann
Turing	ni	$G\ddot{o}del$

is an Armstrong table for Σ and R_s . Indeed, the maximal set families are as before:

- $max_{\Sigma,R_s}(Emp) = \{\{Dept, Mgr\}\},\$
- $max_{\Sigma,R_s}(Dept) = \{\{Mgr\}\}, and$
- $max_{\Sigma,R_s}(Mgr) = \{\{Emp\}\}.$

The agree sets of r are:

- (Ø,Ø),
- $({Dept, Mgr}, {Dept, Mgr}),$
- $({Mgr}, {Emp, Mgr}),$
- $({Emp}, {Emp, Dept})$.

The conditions 1., 2. and 3. of Theorem 4 are satisfied. Note, in particular, that $(Mgr)^*_{\Sigma,R_s} = \{Mgr\}$ which is a proper subset of $w(Mgr) = \{Emp, Mgr\}$.

Finally, we illustrate Theorem 4 for the case when the input relation is not an Armstrong table for the input FD set and input NFS.

Example 16 Let the relation schema R, FD set Σ , and relation r be given as in Example 3, but let the NFS R_s be {Emp,Mgr}. Recall from Example 12 the maximal sets {Dept, Mgr}, {Mgr} and {Emp}. The agree sets of r are:

- $({Dept, Mgr}, {Dept, Mgr}),$
- $({Mgr}, {Mgr}),, and$
- *(*Ø,Ø*)*.

While condition 2. of Theorem 4 is satisfied, conditions 1. and 3. are both violated. In fact, $\{Emp\} \in max_{\Sigma,R_s}(Mgr) - ag^s(r)$, and $Dept \in total(r) - R_s$.

5.3 Characterizing the satisfaction of FD sets

The next theorem generalizes the result from total relations [113] saying that a total relation satisfies a set of FDs precisely when the agree sets of the total relation are closed attribute sets.

Theorem 5 Let R be some relation schema, let Σ be an FD set, and R_s an NFS over R. For all relations r over R it holds that r satisfies Σ and R_s if and only if $R_s \subseteq total(r)$ and $\forall X \in ag^s(r)(X^*_{\Sigma,R_s} \subseteq w(X))$ holds.

Proof Sufficiency. The proof of Theorem 4 shows that r satisfies Σ , if $\forall X \in ag^s(r)(X_{\Sigma}^* \subseteq w(X))$ holds. Furthermore, if $R_s \subseteq total(r)$, then r is R_s -total.

Necessity. If r satisfies R_s , then $R_s \subseteq total(r)$. It remains to show that $\forall X \in ag^s(r)(X_{\Sigma,R_s}^* \subseteq w(X))$ holds, if r satisfies Σ and R_s . Assume that there is some $X \in ag^s(r)$ and there is some $A \in R$ such that $A \in X_{\Sigma,R_s}^* - w(X)$. Since $A \notin w(X)$, there is some $Y \subseteq R$ such that $(X,Y) \in ag(r)$ and $A \notin Y$. Consequently, there are distinct $t, t' \in r$ such that $X = ag^s(t,t')$ and $Y = ag^w(t,t')$. That is, t[X] = t'[X] and t,t' are X-total, and $t[A] \neq t'[A]$. Hence, r violates the FD $X \to A$. From $A \in X_{\Sigma,R_s}^*$ it follows that $X \to A \in \Sigma_{R_s}^*$. The definition of implication shows that r violates Σ .

Example 17 Example 16 shows that the relation from Example 3 satisfies the FD set Σ and NFS R_s from Example 3.

6 Computation of Armstrong Tables

Following Section 4 we need to ask whether there is a reasonable class of FDs in the presence of an NFS that does enjoy Armstrong tables. We can give an affirmative answer: the focus on standard FDs guarantees the existence of Armstrong tables in the presence of an NFS. Instead of showing the mere existence of Armstrong tables, we establish an algorithm that computes an Armstrong table for an arbitrarily given relation schema, an arbitrarily given set of standard FDs and an arbitrarily given NFS over the relation schema.

Our refinement of the notion of an agree set in Definition 1 enables us to utilize a similar idea that Mannila and Räihä used to compute Armstrong relations for sets of FDs over total relations [42]. We specify tuples that have *strong* agree sets on the maximal sets of the attributes. In addition, these tuples are total and unique on the null-free subschema. Before we can present the computation of these Armstrong tables, we establish an algorithm for computing the families of maximal sets.

Lemma 1 Let R be a relation schema, R_s a null-free subschema over R, and $\Sigma = \Sigma' \cup \{X \to A\}$ a set of standard FDs over R. For $WC \subseteq R$, it takes $\mathcal{O}(|R| \times ||\Sigma||)$ time to test whether $W \in \max_{\Sigma, R_s}(C)$.

Proof Using the NFSclosure-algorithm, $C \notin W^*_{\Sigma,R_s}$ can be checked in time $\mathcal{O}(||\Sigma||)$, and $C \in (WB)^*_{\Sigma,R_s}$ for all $B \in R - W$ can be checked in time $\mathcal{O}(|R| \times ||\Sigma||)$.

We use $mtest(W, C, R, R_s, \Sigma)$ to denote the test if $W \in max_{\Sigma,R_s}(C)$ from Lemma 1. The maximal sets for R with respect to Σ and R_s can be computed by testing all subsets of R. This, however, will hardly be efficient. The following result establishes an iterative approach for computing the maximal sets for R with respect to Σ and R_s . The algorithm starts with the maximal sets for R with respect to an empty FD set in the presence of R_s , and then adds the FDs of Σ one by one while monitoring the resulting changes to the family of maximal sets.

Theorem 6 Let R be a relation schema, R_s a null-free subschema over R, and $\Sigma = \Sigma' \cup \{X \to A\}$ a set of standard functional dependencies over R. For $C \in R$ let $V \in max_{\Sigma,R_s}(C)$. Then $V \in max_{\Sigma',R_s}(C)$ or $(C = A \text{ or } A \in R_s)$ holds and there is some $B \in X - V$ such that

- i) $VB \in max_{\Sigma',R_s}(C)$, if $X \not\subseteq R_s$, or
- ii) $V = W \cap Z$ for some $W \in max_{\Sigma',R_s}(C)$ and some $Z \in max_{\Sigma',R_s}(B)$.

The proof uses the following simple observation:

Remark 4 Let $\Sigma = \Sigma' \cup \{X \to A\}$ and $U \subseteq R$. When $A \in U^*_{\Sigma,R_s}$ we have

$$U_{\Sigma,R_s}^* = \begin{cases} (UA)_{\Sigma,R_s}^* = (UA)_{\Sigma',R_s}^* &, \text{ if } A \in R_s \\ U_{\Sigma',R_s}^* A &, \text{ otherwise} \end{cases}$$

while $U_{\Sigma,R_s}^* = U_{\Sigma',R_s}^*$ holds when $A \notin U_{\Sigma,R_s}^*$. Furthermore, the following statements are equivalent: a) $U_{\Sigma',R_s}^* \subset U_{\Sigma,R_s}^*$, b) $A \in U_{\Sigma,R_s}^* - U_{\Sigma',R_s}^*$, and c) $A \notin U_{\Sigma',R_s}^*$ and $X \subseteq U_{R_s} \cap U_{\Sigma',R_s}^*$, cf. Algorithm 1.

Proof From $V \in max_{\Sigma,R_s}(C)$ and $V \subseteq V_{\Sigma',R_s}^* \subseteq V_{\Sigma,R_s}^*$ we get $C \notin V_{\Sigma,R_s}^*$ and $C \notin V_{\Sigma',R_s}^*$. If $V \in max_{\Sigma',R_s}(C)$ we are done. Otherwise, there is some $W \in max_{\Sigma',R_s}(C)$ with $V \subset W$. By $V \in max_{\Sigma,R_s}(C)$ we obtain $C \in W_{\Sigma,R_s}^* - W_{\Sigma',R_s}^*$, so that $W_{\Sigma',R_s}^* \subset W_{\Sigma,R_s}^*$. When applying Remark 4 to the set W, we note $A \in W_{\Sigma,R_s}^* - W_{\Sigma',R_s}^*$ and, therefore, $(C = A \text{ or } A \in R_s)$.

Remark 4 for W further yields $A \notin W_{\Sigma',R_s}^*$ and $X \subseteq WR_s \cap W_{\Sigma',R_s}^*$. Hence, $A \notin V_{\Sigma',R_s}^*$ as $V \subset W$. Assume $A \in V_{\Sigma,R_s}^* - V_{\Sigma',R_s}^*$. Then, $C \neq A$ as $C \notin V_{\Sigma,R_s}^*$, and thus $A \in R_s$. When applying Remark 4 to the set V, we observe $V_{\Sigma,R_s}^* = (VA)_{\Sigma,R_s}^* = (VA)_{\Sigma',R_s}^*$. Hence, $C \notin (VA)_{\Sigma,R_s}^*$. By $V \in max_{\Sigma,R_s}(C)$ we obtain $A \in V$. Thus $A \in V \subseteq W \subseteq W_{\Sigma',R_s}^*$ which contradicts $A \in W_{\Sigma,R_s}^* - W_{\Sigma',R_s}^*$. Hence, our assumption is false, and $A \notin V_{\Sigma,R_s}^*$ holds. From $A \notin V_{\Sigma,R_s}^*$ and Remark 4 for V we conclude $V_{\Sigma,R_s}^* = V_{\Sigma',R_s}^*$, and $X \not\subseteq VR_s \cap V_{\Sigma',R_s}^*$, that is, $(X \not\subseteq VR_s \text{ or } X \not\subseteq V_{\Sigma',R_s}^*)$. By $X \subseteq WR_s$ we obtain $((X - R_s) \cap (W - V) \neq \emptyset$ or $X \not\subseteq V_{\Sigma',R_s}^*)$. Therefore, V is a subset of a member U of

$$\mathcal{V} := \{ W - B : B \in X - R_s \} \cup \\ \{ W \cap Z : W \not\subseteq Z \in \bigcup_{B \in X} max_{\Sigma', R_s}(B) \}.$$

From $U \subset W$ and $W \in max_{\Sigma',R_s}(C)$ we get $C \notin U^*_{\Sigma',R_s}$. By definition of \mathcal{V} , we observe $X \not\subseteq UR_s \cap U^*_{\Sigma',R_s}$ which yields $U^*_{\Sigma,R_s} = U^*_{\Sigma',R_s}$ by Remark 4. Thus, we obtain $C \notin U^*_{\Sigma,R_s}$. From $V \in max_{\Sigma,R_s}(C)$ we derive V = U, that is, V itself is a member of \mathcal{V} . This concludes the proof of Theorem 6. Theorem 6 extends Theorem 4 in [42, page 137] from the special case where $R_s = R$ to an arbitrary NFS. Indeed, for total relations case i) of Theorem 6 cannot occur, and the condition $X \subseteq R_s$ and $(C = A \text{ or } A \in R_s)$ is always satisfied. That is, Theorem 6 becomes Theorem 4 in [42, page 137] for the special case of total relations. Based on Theorem 6 we will now establish an algorithm for computing the maximal set families for a given relation schema R, a given set Σ of standard FDs and a given NFS R_s over R. Recall that $mtest(W, C, R, R_s, \Sigma)$ denotes the test whether $W \in max_{\Sigma,R_s}(C)$ from Lemma 1. Theorem 6 shows that Algorithm 2 on page 28 is correct.

Theorem 7 Algorithm 2, on input (R, Σ, R_s) , computes the families $max_{\Sigma,R_s}(A)$ for every $A \in R$.

The following example illustrates that the evolution of maximal sets is conceptually different from the case of total relations, in particular for the case of a new FD whose LHS is not a subset of R_s .

Example 18 Let R = KLM, $\Sigma' = \{K \to L, L \to M\}$, and $\Sigma = \Sigma' \cup \{M \to K\}$. Assuming that $L \notin R_s$ we find $\max_{\Sigma',R_s}(K) = \{LM\}$, $\max_{\Sigma',R_s}(L) = \{M\}$, and $\max_{\Sigma',R_s}(M) = \{K\}$. The remaining families of maximal sets depend on the precise choice of R_s . If $R_s = \emptyset$ we have $\max_{\Sigma,R_s}(K) = \{L\}$ and $\max_{\Sigma,R_s}(L) = \{M\}$. This is case i) in Theorem 6. It is important to note that $\max_{\Sigma,R_s}(K)$ is not an intersection of maximal sets for R with respect to Σ' and R_s . If $R_s = M$ we have $\max_{\Sigma,R_s}(K) = \emptyset$ and $\max_{\Sigma,R_s}(L) = \{M\}$. This is the case of Theorem 6 when a new dependency is added whose RHS is not declared null-free. Note that only members of $\max_{\Sigma',R_s}(K)$ can change as K is the attribute on the RHS of the new dependency. Finally, if $R_s = KM$ we have $\max_{\Sigma,R_s}(K) = \emptyset$ and $\max_{\Sigma,R_s}(K) = \emptyset$ and $\max_{\Sigma,R_s}(L) = \emptyset$. This is the case of Theorem 6 where a new dependency is added whose RHS is not declared null-free. Only this special case can be regarded as similar to [42, Theorem 4].

Algorithm 3 on page 29 computes an Armstrong table for an arbitrary set Σ of standard FDs and an NFS R_s . In the algorithm, A_0 denotes an arbitrary fixed attribute of the underlying relation schema R. For $A \in R$ let $c_{A,1}, c_{A,2}, \ldots \in dom(A)$ denote distinct elements of the domain. The main construction is achieved between lines 5 and 13. The steps in lines 14-30 guarantee that the output relation is total on exactly those attributes that belong to R_s , and that the relation is subsumption-free. For the soundness proof of Algorithm 3 we utilize Atzeni and Morfuni's axiomatization [46].

Theorem 8 Algorithm 3, on input (R, Σ, R_s) , computes an Armstrong table for Σ and R_s .

Proof Let r denote the output of Algorithm 3. We show first that the output r of Algorithm 3 is a subsumption-free relation. Let $t, t' \in r$ denote two distinct tuples. Suppose t results from some $X \in max(A)$ and t' results from some $Y \in max(A)$. The construction guarantees that $t[A] \neq t'[A]$ and $t[A] \neq ni \neq t'[A]$. Hence, neither of t, t' subsumes the other. Suppose that t results from some $X \in max(A)$ and t' results from some $Y \in max(B)$ where $A \neq B$. If $t'[A] \neq ni$, then $t[A] \neq t'[A]$ and $t[A] \neq t'[A]$

Algorithm 2 Maximal set computation

```
1: procedure MAX-SETS(R, \Sigma, R_s)
 2:
        for all C \in R do
 3:
             max(C) \leftarrow \{R - C\};
        end for
 4:
        \Theta := \emptyset;
 5:
        for all X \to A \in \Sigma do
 6:
            \Theta \leftarrow \Theta \cup \{X \to A\};
 7:
            for all C \in R where (C = A \text{ or } A \in R_s) do
 8:
                nmax(C) \leftarrow max(C);
 9:
                for all W \in max(C) do
10:
11:
                     if not mtest(W, C, R, R_s, \Theta) then
                         nmax(C) \leftarrow nmax(C) - \{W\};
12:
                         for all B \in X do
13:
                             if B \notin R_s and mtest(W - B, C, R, R_s, \Theta) then
14:
                                 nmax(C) \leftarrow nmax(C) \cup \{W - B\};
15:
                             end if
16:
                             for all Z \in max(B) do
17:
                                 if mtest(W \cap Z, C, R, R_s, \Theta) then
18:
                                     nmax(C) \leftarrow nmax(C) \cup \{W \cap Z\};
19:
                                 end if
20:
                             end for
21:
                         end for
22:
                     end if
23:
                end for
24:
            end for
25:
            for all C \in R do
26:
                max(C) \leftarrow nmax(C);
27:
            end for
28:
        end for
29:
        for all c \in R do
30:
             max(C) := \{ W \mid W \in max(C) \text{ and } W \neq \emptyset \};
31:
32:
        end for
        return max(C) for all C \in R;
33:
34: end procedure
```

Algorithm 3 Armstrong table computation

```
1: procedure ARMSTRONG(R, \Sigma, R_s)
              MAX-SETS(R, \Sigma, R_s) (Algorithm 2 on page 28);
 2:
              r := \emptyset;
 3:
              i := 1;
 4:
              for all X \in max(R) do
 5:
                      Z := \{A \in R \mid X \in max(A)\};
 6:
                    Z := \{A \in R \mid A \in R \text{ do} \\ \text{for all } A \in R \text{ do} \\ t_i(A) := \begin{cases} c_{A,i} \text{, if } A \in XZR_s \\ \text{ni} \text{, else} \end{cases}; \\ t_{i+1}(A) := \begin{cases} c_{A,i} \text{, if } A \in X \\ c_{A,i+1} \text{, if } A \in Z(R_s - X) \\ \text{ni} \text{, else} \end{cases};
 7:
 8:
 9:
                     end for
10:
                     r := r \cup \{t_i, t_{i+1}\};
11:
                     i := i + 2;
12:
              end for
13:
              total(r) := \{ A \in R \mid \forall t \in r(t[A] \neq \mathtt{ni}) \};
14:
              if total(r) - R_s \neq \emptyset then
15:
                     if R_s = \emptyset and total(r) = R and |R| > 1 then
16:
                            for all A \in R do

t_i(A) := \begin{cases} c_{A,i} , \text{if } A = A_0 \in R \\ \text{ni} , \text{else} \end{cases};

t_{i+1}(A) := \begin{cases} \text{ni} , \text{if } A = A_0 \in R \\ c_{A,i+1} , \text{else} \end{cases};
17:
18:
19:
                             end for
20:
                             return r := r \cup \{t_i, t_{i+1}\};
21:
                     else
22:
                            for all A \in R do

t_i(A) := \begin{cases} \operatorname{ni} & \text{, if } A \in \mathit{total}(r) - R_s \\ c_{A,i} & \text{, else} \end{cases}
23:
24:
                             end for
25:
                             return r := r \cup \{t_i\}
26:
27:
                     end if
              else
28:
29:
                     return r
              end if
30:
31: end procedure
```

 $ni \neq t'[A]$. If $t[B] \neq ni$, then $t[B] \neq t'[B]$ and $t[B] \neq ni \neq t'[B]$. If t'[A] = ni and t[B] = ni, then neither of t, t' subsumes the other. Finally, if t results from some $X \in max(A)$ and t' results from lines 14-30, or if t and t' result from lines 14-30, then the construction (distinct values) guarantees that neither of t, t' subsumes the other. Hence, r is a subsumption-free relation.

We show r to be an Armstrong table for Σ and R_s .

Let $X \to A \in \Sigma$. Assume that r violates $X \to A$. Then there are distinct $t, t' \in r$ such that t[X] = t'[X] and t, t' are X-total and $t[A] \neq t'[A]$. Since $X \neq \emptyset$, it follows from construction that $\{t, t'\} = \{t_{2i-1}, t_{2i}\}$ for some positive integer i. According to lines 6-11 we conclude that there is some $X' \in max(R)$ and there is some $B \in Z :=$ $\{C \in R \mid X' \in max(C)\}$ such that $X \subseteq X'$ and $A \in ZR_s$. Suppose that $A \in Z$. Consequently, $X \subseteq X' \in max(A)$, i.e., $A \notin (X')_{\Sigma,R_s}^*$. However, due to the soundness of the augmentation rule we conclude that $A \notin X_{\Sigma,R_s}^*$. This means that $X \to A \notin \Sigma_{R_s}^*$ which contradicts $X \to A \in \Sigma$. Suppose now that $A \in R_s - Z$. From $X' \in max(B)$ it follows that $X' \to B \notin \Sigma_{R_s}^*$ and that $X'A \to B \in \Sigma_{R_s}^*$. From $X \to A \in \Sigma_{R_s}^*$ follows $X' \to A \in \Sigma_{R_s}^*$ by the soundness of the augmentation rule. From the soundness of the reflexivity axiom and the union rule we conclude that $X' \to X'A \in \Sigma_{R_s}^*$. An application of the null transitivity rule to $X' \to X'A$, $X'A \to B$ and $A \in R_s$ results in $X' \to B$. Due to the soundness of the null transitivity rule we conclude that $X' \to B \in \Sigma_{R_s}^*$. This contradicts the fact that $X' \in max(B)$. We have just shown that r satisfies Σ . The construction in lines 6-11 ensures that r satisfies the NFS R_s .

It is not difficult to see that the relation r violates all standard FDs $X \to A \notin \Sigma_{R_s}^*$. In fact, by definition of max(A) there is some $X' \subseteq R$ such that $X' \in max(A)$ and $X \subseteq X'$. Lines 7-11 guarantee that there are some distinct $t, t' \in r$ such that t[X'] = t'[X'], t, t' are X'-total and $t[A] \neq t'[A]$. Hence, r violates $X \to A$.

It is now quite easy to see that the relation r is an Armstrong table for Σ and R_s . In fact, lines 14-30 guarantee that r is total on precisely those attributes of R that belong to R_s . Note that $R_s \subseteq total(r)$ always holds due to the construction. Hence, if $total(r) - R_s \neq \emptyset$, then we need to add some tuples with occurrences of **ni** in all columns $A \in total(r) - R_s$. If $R_s = \emptyset$, total(r) = R and |R| > 1, then we require two tuples to ensure that r remains subsumption-free. Otherwise, we can just add a single tuple with occurrences of **ni** in all columns $A \in total(r) - R_s$. Hence, r satisfies precisely those null-free subschema constraints implied by R_s (namely the subsets of R_s).

Example 19 Let EMPLOYMENT = $\{Emp, Dept, Mgr\}$,

 $\Sigma = \{ Emp \to Dept, Dept \to Mgr \}$

and $R_s = \{Emp, Mgr\}$ be as in Example 2. The following table illustrates the evolution of the maximal set families for the attributes of EMPLOYMENT using Algorithm 2 on input (EMPLOYMENT, Σ , R_s) when the first FD considered is $\sigma = Dept \rightarrow Mgr$:

A	$max_{\emptyset,R_s}(A)$	$max_{\{\sigma\},R_s}(A)$	$max_{\Sigma,R_s}(A)$
Emp	$\{\{Dept, Mgr\}\}$	$\{\{Dept, Mgr\}\}$	$\{ \{ Dept, Mgr \} \}$
Dept	$\{ \{ Emp, Mgr \} \}$	$\{\{Emp, Mgr\}\}$	$\{\{Mgr\}\}$
Mgr	$\{ \{ Emp, Dept \} \}$	$\{\{Emp\}\}$	$\{\{Emp\}\}$

Based on the elements of \max_{Σ,R_s} Algorithm 3 would compute the following Armstrong table

Emp	Dept	Mgr
$c_{Emp,1}$	$c_{Dept,1}$	$c_{Mgr,1}$
$c_{Emp,2}$	$c_{Dept,1}$	$c_{Mgr,1}$
$c_{Emp,3}$	$c_{Dept,3}$	$c_{Mgr,3}$
$c_{Emp,4}$	$c_{Dept,4}$	$c_{Mgr,3}$
$c_{Emp,5}$	ni	$c_{Mgr,5}$
$c_{Emp,5}$	ni	$c_{Mgr,6}$

for Σ and R_s .

Finally, we state some consequences of the results we have established in this section.

Corollary 2 The class of standard FDs enjoys Armstrong tables in the presence of an NFS.

Algorithm 3 can also be used to compute Armstrong relations for Lien's class of standard FDs [47]. In that case, we simply set the null-free subschema to $R_s = \emptyset$, but lines 14-30 are unnecessary since we do not need to consider any NFS.

Corollary 3 Algorithm 3, on input (R, Σ, \emptyset) , computes an Armstrong relation with respect to Σ .

Corollary 4 Lien's class of standard FDs enjoys Armstrong relations.

Note that Algorithm 3 can also be used to compute Armstrong relations for the class of standard FDs over total relations.

7 Complexity Results

We show that the gain in generality we have established for the toolbox of Armstrong relations does not result in a loss of efficiency when compared to the special case of total relations. Firstly, the problem of finding an Armstrong table for a given set of standard FDs and a given NFS remains precisely exponential in the size of the input, as was the case for total relations [40]. Nevertheless, our algorithm for computing Armstrong tables remains quite conservative in the sense that the size of the output is at most quadratic in the size of the best output possible, similar to the special case of total relations [42]. Next we examine the most concise way of representing the information inherent in an FD set. We conclude that already for the special case of total relations neither the representation in form of an FD set or the representation in form of an Armstrong table strictly dominates the other. Finally, we show that the problem of deciding whether there is some Codd key with at most k attributes that is implied by a given set of FDs and an NFS with k attributes is NP-complete for the representation either as FD set or an Armstrong table thereof. Hence, the complexity remains the same as in the special case of total relations [40, 114].

7.1 The time-complexity to find Armstrong tables

The user-friendly representation of an FD set and an NFS in form of an Armstrong table comes, in general, at a price. In fact, the number of tuples in a minimum-sized Armstrong table can be exponential in the number of attributes. Due to this result we cannot design an algorithm for generating Armstrong tables in polynomial time in the worst case. The next result shows that the number of attribute sets maximal for Σ and R_s is a lower bound for the number of agree sets found in any Armstrong table for Σ and R_s . A similar result holds for the special case of total relations [40].

Proposition 1 Let Σ be a set of standard FDs, let R_s be some NFS over some relation schema R, and let r be an Armstrong table for Σ and R_s . Then $|max_{\Sigma,R_s}(R)| \leq |ag(r)| \leq {\binom{|r|}{2}}$.

Proof The first condition of Theorem 4 implies that $|max_{\Sigma,R_s}(R)| \leq |ag^s(r)|$. Moreover, $|ag^s(r)| \leq |ag(r)|$, and $|ag(r)| \leq {|r| \choose 2}$ since every distinct pair of distinct tuples in r has precisely one agree set.

We recall what we mean by precisely exponential [40]. Firstly, it means that there is an algorithm for computing an Armstrong table, given a set Σ of standard FDs and an NFS R_s , where the running time of the algorithm is exponential in the number of attributes. Secondly, it means that there is a set Σ of standard FDs and an NFS R_s in which the number of tuples in each minimum-sized Armstrong table for Σ and R_s is exponential — thus, an exponential amount of time is required in this case simply to write down the relation.

Proposition 2 The complexity of finding an Armstrong table, given a set of standard functional dependencies and a null-free subschema, is precisely exponential in the number of attributes.

Proof The time complexity of Algorithm 3 is dominated by that of Algorithm 2 which runs clearly in time exponential in the number of attributes.

It remains to show that there is a set Σ of standard FDs and an NFS R_s for which the number of tuples in each Armstrong table for Σ and R_s is exponential in the number of attributes. According to Proposition 1 it suffices to find a set Σ of standard FDs such that $max_{\Sigma,R_s}(R)$ is exponential in the number of attributes. Such a set Σ is given by

$$\bigcup_{1 \le i \le n} \{\{A_{2i-1}, A_{2i}\} \to B\}$$

and the NFS $R_s = A_1 \cdots A_{2n} B$. This is the same set that Beeri, Dowd, Fagin and Statman used to show that the time complexity of finding an Armstrong relation for FDs over total relations takes at least exponential time in the number of attributes [40]. This set works here for the same purpose since all FDs in Σ have the same right-hand side.

7.2 The size of minimum-sized Armstrong tables

Despite the general worst-case exponential complexity in the number of attributes, Algorithm 3 is a fairly simple algorithm for generating Armstrong tables that is, as we show now, quite conservative in its use of time.

Let the size of an Armstrong table be defined as the number of tuples that it contains. In practice, the most appealing Armstrong table for an FD set Σ should be of minimum size. The reason is that a small number of tuples is easier to comprehend for humans. Therefore, it is a practical question to ask how many tuples a minimum-sized Armstrong table requires. An Armstrong table r for Σ and R_s is said to be *minimum-sized* if there is no Armstrong table r' for Σ and R_s such that |r'| < |r|. That is, for a minimum-sized Armstrong table for Σ and R_s there is no Armstrong table for Σ and R_s with a smaller number of tuples.

Proposition 3 Let Σ be a set of standard FDs, let R_s be some NFS over some relation schema R, and let r be a minimum-sized Armstrong table for Σ and R_s . Then

$$\frac{\sqrt{1+8\cdot|\max_{\Sigma,R_s}(R)|}}{2} \le |r| \le 2 \times |\max_{\Sigma,R_s}(R)| + 1.$$

Proof The lower bound follows from Proposition 1. Indeed, it follows that $|max_{\Sigma,R_s}(R)| \leq {\binom{|r|}{2}}$. Consequently, we have that $\frac{\sqrt{1+8} \cdot |max_{\Sigma,R_s}(R)|}{2} \leq |r|$. The upper bound $2 \times |max_{\Sigma,R_s}(R)| + 2$ follows immediately from Theorem 8. However, Algorithm 3 outputs an Armstrong table of size $2 \times |max_{\Sigma,R_s}(R)| + 2$ if and only if total(r) = R holds before line 14 and $R_s = \emptyset$. We have $B \in total(r)$ before line 14 if and only if $B \in XA \cup R_s$ for every maximal set $X \in max_{\Sigma,R_s}(A)$ and all $A \in R$. Therefore, we have total(r) = R before line 14 and $R_s = \emptyset$ if and only if $max_{\Sigma,R_s}(A) = \{R-A\}$ for all $A \in R$ and $R_s = \emptyset$. This again holds if and only if $\Sigma = \emptyset$ and $R_s = \emptyset$. In this special case and when |R| > 1, one may use the Algorithm from Mannila and Räihä [42] to compute an Armstrong relation for Σ of size $max_{\Sigma,R_s}(R) + 1 = |R| + 1$ and add two tuples according to lines 14-30 of Algorithm 3 to obtain an Armstrong table for Σ and R_s . If |R| = 1, then the singleton relation consisting of the tuple $t = \mathbf{ni}$ is an Armstrong table for Σ and R_s .

We conclude that Algorithm 3 always computes an Armstrong table of reasonably small size.

Corollary 5 On input (R, Σ, R_s) , Algorithm 3 computes an Armstrong table for Σ and R_s whose size is at most quadratic in the size of a minimum-sized Armstrong table for Σ and R_s .

There are also instances for which Algorithm 1 computes a minimum-sized Armstrong table. For each $n \geq 3$, such an instance is given by $R_n = A_1 \cdots A_n$, $R_s = \emptyset$, and $\Sigma = \bigcup_{i=1}^{n-1} \{A_i \to A_{i+1}\} \cup \{A_n \to A_1\}.$

Theorem 9 There is some relation schema R, some NFS R_s and some standard FD set Σ over R such that Algorithm 3 computes a minimum-sized Armstrong table for Σ and R_s .

Proof Let $R = A_1 \cdots A_n$ with $n \geq 3$, $R_s = \emptyset$, and Σ consist of the FDs $A_i \to A_{i+1}$ with $i = 1, \ldots, n$. We find $max_{\Sigma,R_s}(A_i) = \{R - A_{i-1}A_i\}$ for $i = 1, \ldots, n$ (with the convention that $A_i = A_{i+n} = A_{i-n}$). Let r be any Armstrong table for Σ and R_s . We show that its size it at least $2 \times |max_{\Sigma,R_s}(R)|$. By Theorem 4 there are tuples $t_i, t'_i \in r$ for all $i = 1, \ldots, n$ such that $ag^s(t_i, t'_i) = R - A_{i-1}A_i$ and $t_i[A_i] \neq t'_i[A_i]$. We conclude $t_i[A_{i-1}] = ni = t'_i[A_{i-1}]$ for $i = 1, \ldots, n$ since r satisfies the FD $A_{i-2} \to A_{i-1}$, and also $t_i[A_i] \neq ni \neq t'_i[A_i]$ for $i = 1, \ldots, n$ since r is subsumption-free. Hence, $total(\{t_i\}) =$ $R - A_{i-1} = total(\{t'_i\}) = R - A_{i-1}$ for $i = 1, \ldots, n$. As $t_1, t'_1, \ldots, t_n, t'_n$ are mutually distinct, r has size at least $2 \times n = 2 \times |max_{\Sigma,R_s}(R)|$.

7.3 The size of representations

We show that, in general, there is no most concise way of representing the information inherent in a set of standard FDs and a null-free subschema. We have already seen a case where the representation using Armstrong tables can be exponentially larger than the best equivalent FD set, see Proposition 2.

Corollary 6 There is some standard FD set Σ and an NFS R_s such that Σ has size $\mathcal{O}(n)$, and the size of a minimum-sized Armstrong table for Σ and R_s is $\mathcal{O}(2^{n/2})$.

The following theorem shows that in other cases, the representation using Armstrong tables can be exponentially smaller than the best representation using FD sets. Extending Maier's notion of an optimal cover from total relations [115], for an FD set Σ and an NFS R_s we call an FD set Σ' an *optimal cover* of Σ with respect to R_s if

- Σ' is a cover of Σ with respect to R_s , i.e., for every FD $\sigma \in \Sigma$ we have $\Sigma' \models_{R_s} \sigma$; and for every FD $\sigma' \in \Sigma'$ we have $\Sigma \models_{R_s} \sigma'$; and
- there is no cover Σ'' of Σ with respect to R_s such that Σ'' contains fewer symbol occurrences than Σ' (repeated symbol occurrences are counted as many times as they occur).

Let $R = A_1 B_1 \cdots A_n B_n C$, $R_s = R$ and

$$\Sigma = \{X_1, \dots, X_n \to C \mid \forall i = 1, \dots, n(X_i \in \{A_i, B_i\})\}.$$

Then Σ is its own optimal cover with respect to R_s , but there is an Armstrong table for Σ and R_s where the number of tuples is in $\mathcal{O}(n)$.

Theorem 10 There is some relation schema R, some NFS R_s and some standard FD set Σ over R such that there is an Armstrong table for Σ and R_s where the number of tuples is in $\mathcal{O}(n)$, and the optimal cover of Σ with respect to R_s has size $\mathcal{O}(2^n)$.

Proof Let $R = A_1 B_1 \cdots A_n B_n C$, $R_s = R$ and

$$\Sigma = \{X_1, \dots, X_n \to C \mid \forall i = 1, \dots, n(X_i \in \{A_i, B_i\})\}$$

We show that Σ is the optimal cover of Σ with respect to R_s , but there is an Armstrong table for Σ and R_s where the number of tuples is in $\mathcal{O}(n)$.

We will show first that Σ is non-redundant (no subset of Σ implies all FDs in Σ), and then show that Σ is an optimal cover of itself. We note that for every FD $\sigma \in \Sigma$, where $X = LHS(\sigma)$ denotes the attribute set on the left-hand side of σ , the closure $X^*_{(\Sigma - \{\sigma\}, R_s)}$ of X with respect to $\Sigma - \{\sigma\}$ and R_s is X itself, i.e., $X^*_{(\Sigma - \{\sigma\}, R_s)} = X$. The reason is that there is no $\sigma' \in \Sigma - \{\sigma\}$ such that $LHS(\sigma') \subseteq X$. Hence, $C \notin X^*_{(\Sigma - \{\sigma\}, R_s)}$ and we conclude that σ is not implied by $\Sigma - \{\sigma\}$ and R_s . That is, Σ is non-redundant. Next we remark that every optimal cover Σ' of Σ with respect to R_s contains only FDs $X \to Y$ such that Y = C. Suppose, to the contrary, that there is some FD $X \to Y$ in Σ' such that $Y \neq C$. If Y - X = C and $Y \cap X \neq \emptyset$, then Σ' is not optimal since $(\Sigma' - \{X \to Y\}) \cup \{X \to Y - X\}$ is equivalent to Σ but contains less attributes than Σ' . If $Y - X = \emptyset$, then $\Sigma' - \{X \to Y\}$ is equivalent to Σ but contains less symbol occurrences than Σ' . If $Y - X \neq \emptyset$ and $Y - X \neq C$, then $\Sigma \not\models_{R_s} X \to Y$ and, therefore, Σ' is not a cover of Σ with respect to R_s . Moreover, every FD $X \to Y$ in an optimal cover Σ' of Σ with respect to R_s satisfies that $C \notin X$. If there was an FD $X \to C \in \Sigma'$ and $C \in X$, then

$$(\Sigma' - \{X \to C\}) \cup \{X - C \to C\}$$

is equivalent to Σ but contains less attributes than Σ' .

Next we prove that there is no cover Σ' of Σ with respect to R_s with a smaller number of attribute occurrences. Suppose there were an optimal cover Σ' of Σ with respect to R_s with a fewer number of attribute occurrences than Σ . Then for all σ' in Σ' it is the case that $\Sigma \models_{R_s} \sigma'$. Consequently, there must be some $\sigma \in \Sigma$ such that $LHS(\sigma) \subseteq LHS(\sigma')$. Suppose every FD $\sigma \in \Sigma$ has the property that $LHS(\sigma) \subseteq LHS(\sigma')$ for a different FD $\sigma' \in \Sigma'$. Then Σ' contains at least as many attribute occurrences as Σ , a contradiction. Otherwise, there is a proper subset Σ'' of Σ such that every FD $\sigma' \in \Sigma'$ has the property that $LHS(\sigma) \subseteq LHS(\sigma')$ for some $\sigma \in \Sigma''$. Consequently, Σ'' implies every FD in Σ' with respect to R_s and therefore also every FD in Σ . This, however, is impossible since Σ is non-redundant.

Thus we have just shown that Σ is its own optimal cover with respect to R_s , and thus exponential in the number of attributes. Now we show that there is an Armstrong table for Σ and R_s where the number of tuples is in $\mathcal{O}(n)$. It suffices to show that the set $max_{\Sigma,R_s}(R)$ contains a number of elements that is linear in the number of attributes. For each $i = 1, \ldots, n$ we have $max_{\Sigma,R_s}(A_i) = R - A_i$, and $max_{\Sigma,R_s}(B_i) = R - B_i$. These are 2n different maximal sets in total. The set $max_{\Sigma,R_s}(C)$ consists of the following n elements: $R - A_i B_i C, i = 1, \ldots, n$. Therefore, $max_{\Sigma,R_s}(R)$ has 3n different elements. Using Mannila and Räihä's algorithm [42] (which applies since we are in the case where $R_s = R$) we can easily create an Armstrong table for Σ and R_s that has 3n + 1 tuples only.

We can see that the representation in form of an Armstrong table can offer tremendous space savings over the representation as an FD set, and vice versa.

7.4 The time-complexity to find a Codd key

Finally, it follows easily that the well-known problem [40, 114] of deciding whether there is a key of size at most k is also NP-complete when functional dependencies and null-free subschemata are given. An input to the problem can be either a relation schema R, a set Σ of FDs and an NFS R_s over R, and a non-negative integer $k \leq |R_s|$, or a relation schema R, an Armstrong table r over R for some FD set Σ and some NFS R_s , and a non-negative integer $k \leq |R_s|$. For the first type of input the problem is to decide whether there is a Codd key Codd(X) with $|X| \leq k$ such that $\Sigma \models_{R_s} Codd(X)$. For the second type of input the problem is to decide whether there is a Codd key Codd(X) with $|X| \leq k$ such that r satisfies Codd(X). We say that we are deciding the problem if there is a Codd key of size at most k for either type of input. For the hardness part of the next result we simply choose R_s to be R [40].

Proposition 4 For either type of input, the problem of deciding whether there is a Codd key of size at most k is NP-complete.

8 Extensions of Our Results

In this section we establish some extensions of our results. We analyze the impact of lifting the restriction to subsumption-free relations, and study Armstrong tables for standard functional dependencies with respect to the class of standard and non-standard functional dependencies.

8.1 Subsumption in SQL tables

So far, we have imposed the following restriction: No relation in the database shall contain two tuples t_1 and t_2 such that t_1 subsumes t_2 . The reason was to keep consistency with previous work in this area [46, 47, 112, 109]. In this subsection we analyze the impact on the validity of our results when lifting the restriction.

In general, all the results on Armstrong tables we have established for the class of functional dependencies remain valid. Furthermore, if a Codd key Codd(X) has been specified on the schema, then all the relations over the schema will be subsumption-free. For the remainder of the discussion we therefore look at schemata where no Codd key has been specified. In this case, Theorem 1 fails. In fact, there are now relations over R that satisfy the FD $X \to R$, but still violate the uniqueness constraint unique(X). As a consequence, there is a real interest in studying the combined class of uniqueness constraints and functional dependencies in the presence of an NFS. We exclude the nonstandard uniqueness constraint $unique(\emptyset)$ from our investigation since it would force our relations to contain at most one tuple. Note that for a set Σ of uniqueness constraints and FDs, and an NFS R_s over R, the following holds: $\Sigma \models_{R_s} Codd(X)$ if and only if $\Sigma \models_{R_s} unique(X)$ and $X \subseteq R_s$. Hence, over relations in which subsumption is allowed the combined class of uniqueness constraints and FDs in the presence of an NFS subsumes the class of Codd keys. For the remainder of this discussion we analyze how to compute Armstrong tables for the combined class of uniqueness constraints and FDs in the presence of an NFS. Given a set $\Sigma = \Sigma_{\text{unique}} \cup \Sigma_{\text{FD}}$ of standard uniqueness constraints and standard FDs over relation schema R, we first use Algorithm 3, exclusive of lines 14-30, to compute an Armstrong table r' for the set $\Sigma_{\text{unique}}^{\text{FD}} \cup \Sigma_{\text{FD}}$ of standard FDs. Here,

$$\Sigma_{\text{unique}}^{\text{FD}} := \{ X \to R \mid unique(X) \in \Sigma_{\text{unique}} \}$$

consists of the FDs implied by the uniqueness constraints in Σ_{unique} . However, we need to add further tuples to r' in order to obtain an Armstrong table for Σ . In fact, we need to ensure additionally that all uniqueness constraints unique(X) over R are violated that are not implied by Σ . If Σ does not imply the FD $X \to R$, then r' already violates unique(X). Therefore, it suffices to look at the following family of attribute sets:

$$\begin{array}{lll} dup_{\Sigma,R_s}(R) &=& \{X \subseteq R \mid \Sigma \models_{R_s} X \to R \land \\ & \Sigma \not\models_{R_s} unique(X) \land \\ & \forall A \in R - X(\Sigma \models_{R_s} unique(XA))\}. \end{array}$$

To compute $dup_{\Sigma,R_s}(R)$ we generate the hyper-graph $\mathcal{H} = (V, E)$ with vertex set V = Rand the set

 $E = \{K - R_s \mid unique(K) \in \Sigma_{unique}\}$

as hyper-edges. From this we obtain $dup_{\Sigma,R_s}(R)$ as

$$dup_{\Sigma,R_s}(R) = \{R - X \mid X \in Tr(\mathcal{H}) \land \\ \forall M \in max_{\Sigma,R_s}(R)(R - X \not\subseteq M)\}.$$

where $Tr(\mathcal{H})$ denotes the minimal transversals of the hyper-graph \mathcal{H} [70]. Algorithm 4 on page 38 extends Algorithm 3 on page 29 in order to compute an Armstrong table for a set of standard uniqueness constraints and FDs in the presence of an NFS, considering that relations are not necessarily subsumption-free. In particular, lines 3-11 add to the table tuples that strongly agree on the attribute sets in $dup_{\Sigma,R_s}(R)$. This ensures that the resulting table violates all uniqueness constraints not implied by Σ . Note that for all $X \in dup_{\Sigma,R_s}(R)$ we have $R_s \subseteq X$. Lines 12-20 simplify lines 14-30 from Algorithm 3 since we do not need to ensure that the resulting Armstrong table is subsumption-free.

Example 20 Consider again the relation schema EMPLOYMENT with attribute set Emp, Dept and Mgr, FD set Σ_{FD} that consists of $Emp \rightarrow Dept$ and $Dept \rightarrow Mgr$, and $NFS R_s$ consisting of Emp and Mgr. In particular, let Σ_{unique} be empty. Then the relation

Emp	Dept	Mgr
Hilbert	Math	Gauss
Hilbert	Math	Gauss
Py thag or as	Math	Gauss
Einstein	Physics	Gauss
Turing	ni	von Neumann
Turing	ni	$G\ddot{o}del$

Algorithm 4 Armstrong table computation revised

```
1: procedure Armstrong-revised(\Sigma, T_s, T)
           Let r be the output of Algorithm 3 from lines 2-13 on (R, \Sigma_{\text{unique}}^{\text{FD}} \cup \Sigma_{\text{FD}}, R_s);
 2:
           if \neg \exists X(unique(X) \in \Sigma_{unique} \land X \subseteq R_s) then
 3:
                 for all X \in dup_{\Sigma,R_s}(R) do
 4:
                      for all A \in \overline{R} do

t_i(A) = t_{i+1}(A) := \begin{cases} c_{A,i} , \text{ if } A \in X \\ \texttt{ni} , \text{ else} \end{cases};
 5:
 6:
 7:
                       end for
                      r := r \cup \{t_i, t_{i+1}\};
 8:
                      i := i + 2;
 9:
                 end for
10:
           end if
11:
           total(r) := \{ A \in R \mid \forall t \in r(t[A] \neq \texttt{ni}) \};
12:
           if total(r) - R_s \neq \emptyset then
13:
                for all A \in R do

t_i(A) := \begin{cases} \text{ni} , \text{if } A \in total(r) - R_s \\ c_{A,i}, \text{else} \end{cases}
14:
15:
                 end for
16:
                 return r := r \cup \{t_i\}
17:
18:
           else
19:
                 return r
           end if
20:
21: end procedure
```

is an Armstrong table for $\Sigma_{unique} \cup \Sigma_{FD}$ in the presence of R_s . Here, $dup_{\Sigma,R_s}(R) = \{R\}$ and the first two tuples are duplicates illustrating that no uniqueness constraint is considered meaningful with the present design.

We now consider Beeri and Bernstein's famous example that was originally used to show that dependency-preserving Boyce-Codd Normal Form decompositions cannot always be obtained in the relational model of data [56]. Here, we apply our toolbox of Armstrong tables to the relational approximation of the application domain to derive a concise SQL table definition.

Example 21 Suppose a team of data engineers has decided to use the three attributes Address, City and ZIP as part of a contact address management system. They have identified the set Σ_{FD} with Address, City \rightarrow ZIP and ZIP \rightarrow City as a first approximation for capturing the semantics of the underlying application domain. (This is the example from [56].) Suppose that Σ_{unique} and the NFS R_s remain empty. Using our toolbox the team produces

Address	City	ZIP
03 Hudson St	Manhattan	10001
03 Hudson St	Manhattan	10001
$70 \ King \ St$	Manhattan	10001
$70 \ King \ St$	San Francisco	94107
ni	San Francisco	94129
15 Maxwell St	ni	ni

as an Armstrong table for $\Sigma_{unique} \cup \Sigma_{FD}$ in the presence of R_s . The domain experts inspect the table. They note from the first two tuples that the FD Address, $City \rightarrow ZIP$ should be replaced by the stronger uniqueness constraint unique(Address, City). Hence, the revised constraint set becomes: $\Sigma_{unique} = \{unique(Address, City)\}, \Sigma_{FD} = \{ZIP \rightarrow City\}, and R_s$ is the empty set. For these requirements, they produce

Address	City	ZIP
03 Hudson St	Manhattan	10001
$70 \ King \ St$	Manhattan	10001
$70 \ King \ St$	San Francisco	94107
ni	San Francisco	ni
15 Maxwell St	ni	60609
15 Maxwell St	ni	60609

as an Armstrong table for $\Sigma_{unique} \cup \Sigma_{FD}$ in the presence of R_s . In particular, $dup_{\Sigma,R_s}(R)$ contains the set {Address, ZIP}. The domain experts are quite happy with the table, but are confused about the last two rows. After some discussion there is consensus that there is no reason to allow different rows with the same total values on Address and ZIP. Indeed, the data engineers realize that they have not captured the uniqueness constraint unique(Address, ZIP). (It is important to note here that in the relational model the FD ZIP \rightarrow City implies the key {Address, ZIP}, but over SQL tables the two constraints unique(Address, City) and ZIP \rightarrow City do not imply unique(Address, ZIP).) For this reason, the team revises the set of constraints again: Σ_{unique} consist of unique(Address, City) and unique(Address, ZIP), and $\Sigma_{FD} = \{ZIP \rightarrow City\}$. R_s still remains empty. Then they produce

Address	City	ZIP
03 Hudson St	Manhattan	10001
$70 \ King \ St$	Manhattan	10001
$70 \ King \ St$	San Francisco	94107
ni	San Francisco	ni
15 Maxwell St	ni	60609

as an Armstrong table for $\Sigma_{unique} \cup \Sigma_{FD}$ in the presence of R_s . Here, $dup_{\Sigma,R_s}(R) = \emptyset$. The first two tuples ensure that the uniqueness constraint unique(City, ZIP) is violated. The domain experts note now that there is a need to always have precise information in the Address and ZIP columns. As a consequence, the data engineers decide to include both column headers in the null-free subschema R_s . This is the same as having the following constraints: Codd(Address,ZIP), unique(Address,City) and FD ZIP \rightarrow City. The table

Address	City	ZIP
03 Hudson St	Manhattan	10001
$70 \ King \ St$	Manhattan	10001
$70 \ King \ St$	San Francisco	94107
35 Lincoln Blvd	San Francisco	94129
15 Maxwell St	ni	60609

is an Armstrong table for this design choice. One possible SQL table definition is the following:

CREATE TABLE CONTACT (Address VARCHAR, City VARCHAR, ZIP INT, UNIQUE(Address, City), PRIMARY KEY(Address, ZIP), CHECK(Q = 0));

The state assertion is based on the following query Q:

SELECT COUNT(*) FROM CONTACT c1 WHERE c1.ZIP IN (SELECT ZIP FROM CONTACT c2 WHERE c1.ZIP=c2.ZIP AND (c1.City <> c2.City OR (c1.City IS NULL AND c2.City IS NOT NULL) OR (c1.City IS NOT NULL AND c2.City IS NULL)));

and can be enforced on the data or middle tier.

8.2 Armstrong tables and non-standard FDs

One may also investigate the properties of Armstrong tables for a set of standard FDs and an NFS with respect to the class of all FDs. More formally, a relation r is said to be an Armstrong table for a set Σ of standard FDs and an NFS R_s in the world of all FDs, if $total(r) = R_s$ and the following holds for all standard and non-standard FDs σ : r satisfies σ if and only if $\Sigma \models_{R_s} \sigma$. Since any set of standard FDs does not imply any non-trivial non-standard FD, this definition can be simplified as follows. A relation r is an Armstrong table for a set Σ of standard FDs and an NFS R_s in the world of all FDs, if $total(r) = R_s$, r violates every non-trivial non-standard FD, and the following holds for all standard FDs σ : r satisfies σ if and only if $\Sigma \models_{R_s} \sigma$.

Every Armstrong table r for a set Σ of standard FDs and an NFS R_s can be easily modified to obtain an Armstrong table for Σ and R_s in the world of all FDs: just add to ran additional tuple with previously unused domain values different from ni. An exception is the special case where the underlying relation schema consists of only one attribute A and the NFS $R_s = \emptyset$. In this case, no subsumption-free relation can simultaneously violate the NFS R_s and the non-trivial, non-standard FD $\emptyset \to A$. Consequently, no Armstrong table exists for the empty FD set Σ and the NFS $R_s = \emptyset$ in the world of all FDs, if the underlying relation schema consists of only one attribute. Note, however, that the singleton consisting of the tuple t = ni is an Armstrong table for the empty FD set Σ and the NFS $R_s = \emptyset$ in the "world of all standard FDs". Hence, there are relations that are Armstrong tables in the world of all standard FDs, but not an Armstrong table in the world of all FDs. Note that every relation that is an Armstrong table in the world of all FDs is also an Armstrong table in the world of all standard FDs.

How do the results for standard FD sets from the previous sections change in the world of all FDs? Regarding the notion of maximal sets in Definition 3 it is intuitive to allow also empty attribute sets to be maximal for any attribute. The characterization of Armstrong tables in Theorem 3 carries over to the world of all FDs, if the first condition is true for all attribute sets X including the empty one. The characterization of Theorem 4 carries over to the world of all FDs as it is. Theorem 5 carries over with no change since it is not concerned with the violation of any non-standard FDs. The next example illustrates these points.

Example 22 Consider $R = \{Dept, Mgr\}$ with standard FD set $\Sigma = \{Dept \rightarrow Mgr\}$ and NFS $R_s = R$. The relation r

is an Armstrong table for Σ and R_s , but not an Armstrong table for Σ and R_s in the world of all FDs: the non-standard FD $\emptyset \to Mgr$ is not implied by Σ and R_s , but satisfied by r. Indeed, the empty attribute set X violates the first condition of Theorem 3, and the first condition of Theorem 4 is violated since the empty attribute set is an element of $max_{\Sigma,R_s}(Mgr) - ag^s(r)$. Adding t = (CS, von Neumann) to r results in an Armstrong table for Σ and R_s in the world of all FDs. For the computation of the maximal sets, the only change in Algorithm 2 is concerned with line 31 where we simply remove the condition that W must not be the empty set. Regarding the computation of Armstrong tables, Algorithm 3 does not require any changes. However, the special case where |R| = 1 and $R_s = \emptyset$ should be excluded from the set of all possible inputs to the algorithm since no Armstrong table exists in the world of all FDs, as mentioned before. Corollaries 3 and 4, as well as the results from Section 7 carry over to the world of all FDs.

9 Impact and Applications

In this section we demonstrate the potential impact of our results on various database applications.

An Example Domain. Let us assume that in developing an information system for some manufacturer of electrical goods we identify the processing of orders by retail sellers as a domain of interest. In particular, we define a relation schema ORDER that consists of the attributes $Order \sharp$, $Product \sharp$, Description, Qty, and Total. These show for an order (identified by its order number $Order \sharp$), a product in that order (identified by its unique product number $Product \sharp$), a description Description of that product, the quantity Qty of that product in that order, and the total value Total (in some fixed currency) of that product in that order.

The data engineers of our information system have identified the following FD set Σ to be meaningful: $Order \sharp$, $Product \sharp \rightarrow Qty$, $Product \sharp \rightarrow Description$, and $Product \sharp$, $Qty \rightarrow Total$. The team of data engineers agrees on the NFS $ORDER_s = \{Order \sharp, Product \sharp\}$.

Before the implementation the team validates their design using the Armstrong table for Σ and ORDER_s from Table 2. The engineers observe that the product with *Product* 612 occurs with the same total value *Total* of 25000, but with the different quantities Qty of 75 and 100. After consultation with the domain experts, the team responds by including the additional meaningful FD *Product*, *Total* $\rightarrow Qty$ in Σ . The team further notices that $Codd(Order \sharp, Product \sharp)$ is not implied by their design choice: the product with $Product \sharp$ 834 in the order with $Order \sharp$ 43056 has different total values of 35000 and 40000. After some discussion the team responds by re-defining the NFS to $ORDER_s = \{Order \sharp, Product \sharp, Qty\}$. The inspection of the Armstrong table has resulted in the recognition of additional meaningful business rules.

Efficient Processing of Updates. The main driver of database normalization is to avoid data redundancy to guarantee efficient updates. The Boyce-Codd Normal Form (BCNF) enforces a syntactic condition on relation schemata that eliminates data redundancy in terms of FDs [60, 116, 117]. We assume familiarity with the definition of such terms as BCNF, lossless and dependency-preserving decomposition [2]. Our schema ORDER is not in BCNF with respect to the FD set Σ containing

- Order, Product, $\forall \to Qty$,
- $Product \Rightarrow Description$,
- Product, $Qty \rightarrow Total$, and

$Order \sharp$	$Product \sharp$	Description	Qty	Total
10723	389	Microwave	50	25000
21834	389	Microwave	50	25000
21834	490	Microwave	50	25000
21834	501	Fridge	50	25000
21834	612	Fridge	75	25000
32945	612	Fridge	100	25000
32945	723	Fridge	100	30000
43056	834	Oven	ni	35000
43056	834	Oven	ni	40000
54167	945	ni	200	ni

Table 2: Armstrong table for Σ and $ORDER_s$

• Product, $Total \rightarrow Qty$,

and the NFS $ORDER_s = \{Order \sharp, Product \sharp, Qty\}$. Following standard techniques [2], we decompose ORDER into $PRODUCT = \{Product \sharp, Description\}$ with the FD set $\Sigma_1 = \{Product \sharp \rightarrow Description\}$ and $PRODUCT_s = \{Product \sharp\}, TOTAL = \{Product \sharp, Qty, Total\}$ with the FD set Σ_2 containing

Product, $Qty \rightarrow Total$ and Product, $Total \rightarrow Qty$,

and TOTAL_s = {*Product*#, *Qty*}, and QTY={*Order*#, *Product*#, *Qty*} with FD set $\Sigma_3 =$ {*Order*#, *Product*# $\rightarrow Qty$ }, and QTY_s = {*Order*#, *Product*#, *Qty*}. The three relation schemata represent a lossless and dependency-preserving BCNF decomposition of OR-DER for Σ and ORDER_s. Note that schema ORDER with the original FD set Σ and NFS ORDER_s = {*Order*#, *Product*#} every BCNF decomposition is lossly or not dependencypreserving. Hence, an inspection of the Armstrong table enabled our data engineers to find a database layout that represents all of the business rules, permits efficient consistency checking, is free from data redundancies and update anomalies.

Efficient Processing of Queries. Besides updates, the efficient processing of database queries is also a significant task of DBMSs. Indeed, there has been a lot of research regarding the use of FDs and other constraints in semantic query optimization and query processing with SQL databases including [118, 34, 35, 119, 120, 33, 36, 121, 122]. We will illustrate how the use of our Armstrong tables can result in semantically optimized query re-writings. Recall that the inspection of the Armstrong table enabled the data engineers to identify the meaningful FD *Product*, *Total* $\rightarrow Qty$ and to specify Qty as NOT NULL. The Codd key *Codd*(*Order*, *Product*) is implied by Σ and ORDER_s, in particular. Consider the query that retrieves all combinations of order numbers and quantities associated with the same product and total value. A naive implementation of this query is

SELECT	ORDER'.Order#, ORDER.Qty
FROM	Order, Order AS Order'
WHERE	ORDER.Product =ORDER'.Product AND
	Order.Total#=Order'.Total#

However, since Product, $Total \rightarrow Qty$ has been correctly specified as a meaningful FD, the quantity is the same for each given combination of product and total value. Hence, the query above can be rewritten into

SELECT ORDER.Order#, ORDER.Qty FROM ORDER

which is much simpler and requires no join.

Inference Control is a security mechanism to ensure confidentiality in databases [37, 123, 124]. The objective is to avoid inferences of secrets by users based on their query history and knowledge about the database. Suppose the fact that there is an order of a *Fridge* with a total value of 100000 is a business secret. Hence, the fact that the sentence $\Psi = \exists X_O \exists X_P \exists X_Q \text{ORDER}(X_O, X_P, Fridge, X_Q, 100000)$ is true in the database must not be revealed to unauthorized users. A user may issue the queries:

- $\Phi_1 = \exists X_Q \exists X_T \text{ORDER}(73956, 1645, Fridge, X_Q, X_T)$ and
- $\Phi_2 = \exists X_D \exists X_T \text{ORDER}(73956, 1645, X_D, X_Q, 100000)$

and learn that both queries are true in the current instance, since neither Φ_1 nor Φ_2 individually reveal Ψ . The user may anticipate that the instance satisfies $Order \sharp$, $Product \sharp \rightarrow Description$, Total. Hence, the user may apply this FD to the answers to Φ_1 and Φ_2 to infer that $\exists X_Q \text{ORDER}(73956, 1645, Fridge, X_Q, 100000)$ is true in the database instance. This, however, would reveal the potential secret to the user. Fortunately, the data engineers were able to utilize Armstrong tables to recognize that the FD $Order \sharp$, $Product \ddagger \rightarrow Description$, Total is indeed a meaningful constraint for the application domain. Therefore, the security officer is able to anticipate such an inference, and distort the answers to the queries suitably. In this example, it would suffice to refuse an answer to query Φ_2 after the user has learned Φ_1 . Consequently, Armstrong tables also provide an aid that can help security officers to better understand the opportunities of potential inference attacks on future database instances, and therefore may prove invaluable in preventing such attacks successfully.

Design approaches. Our toolbox of Armstrong tables *complements* any existing approaches to schema design, such as Entity-Relationship modeling [4, 9] and relational normalization [2, 53]. As illustrated by Example 21, Armstrong tables provide data engineers with concise test data that helps them to consolidate their final design choice. It also helps them to communicate and justify their choice to other stakeholders of the database. It is our assumption that the inspection of Armstrong tables exposes those semantically meaningful uniqueness constraints and FDs that are not implied by the set of constraints presently considered meaningful by the data engineers. Recent studies confirm this assumption empirically, already for relations where no duplicate rows and no partial data are allowed to occur [44]. Suppose that the schema {Address, City, ZIP} with Codd(Address, ZIP) and unique(Address, City) constitutes the current design choice. An Armstrong table for the two uniqueness constraints will expose the FD $ZIP \rightarrow City$ as a meaningful constraint that is not captured by this choice. Finally, de-normalization is often applied in practice to increase the efficiency of query processing [2]. Armstrong tables will be particularly useful in identifying the FDs exhibited by de-normalized schemata.

10 Conclusion and Future Work

We have investigated the existence and properties of Armstrong tables for Atzeni and Morfuni's class of FDs and NFS over relations that can contain occurrences of Zaniolo's no information null value. This covers the class of FDs defined over SQL table definitions that contain NOT NULL constraints, and has therefore important implications for the processing of data in real database systems. In contrast to the special case of total relations, we have shown that FDs do not enjoy Armstrong tables, in general. However, the class of standard FDs and NFSs does enjoy Armstrong tables. We have given sufficient and necessary conditions for a given table to be an Armstrong table for a given set of standard FDs and a given NFS. In general, the problem of finding an Armstrong table remains precisely exponential in the number of attributes. However, we have established a provably-correct algorithm that computes an Armstrong table in time that is at most quadratic in the size of a minimum-sized Armstrong table. We have demonstrated that data engineers should utilize not only an abstract representation in form of FD sets, but also the representation in form of an Armstrong table. For SQL tables that are not necessarily subsumption-free we have extended our results to the combined class of uniqueness constraints and FDs in the presence of an NFS. In particular, our Armstrong tables can visualize the delicate interactions of these constraints that can occur in real SQL tables, but not in the relational model of data. Finally, we have illustrated that the utilization of Armstrong tables can lead to the identification of meaningful constraints, and result in database designs that facilitate efficient updates and queries, and are less prone to inference attacks. In summary, our contributions extend well-known results from total relations to SQL tables. Hence, the resulting toolbox can be applied to instances that occur in real database systems, with only small space and no time penalties when compared to the previously studied special case of total relations.

For future work, it would be interesting to investigate the properties of Armstrong tables for the combined class of functional and inclusion dependencies in the presence of a null-free subschema. Levene and Loizou have investigated the implication problem for the class of inclusion dependencies where all attributes are assumed to be NULL [125]. Mannila and Räihä have also developed algorithms to compute Armstrong databases for the class of functional and acyclic inclusion dependencies over total relations [126].

Fagin has shown that Armstrong relations exist for so-called *implicational dependencies* [90]. It would be a worthwhile endeavor to identify broad classes of dependencies that enjoy Armstrong relations in the presence of null values, e.g. multivalued dependencies [48, 47, 112].

A main observation for the areas of impact from Section 9 is that most of the existing theory does not apply to SQL tables. These areas include normalization [127], semantic query optimization [34, 35], consistent query answering [84] and controlled query evaluation [63].

We plan to implement our concepts and algorithms to extend design aids available for total relations [41, 26, 42, 93]. It appears to be intuitive that data engineers find it more difficult to understand the interaction of uniqueness constraints and FDs in the presence of an NFS than that of just FDs over total relations. Hence, Armstrong tables might be of even bigger value than reported for the special case of total relations [44]. In [7] Jagadish et al. develop a user-friendly approach towards spreadsheet design. Here, functional dependencies are discovered from the given spreadsheet data. In other words, the spreadsheet data is an Armstrong relation for the set of FDs discovered. The FDs are then used to recommend auto-completions for updates, and to warn users about potential data entry errors. Based on the users' reply to these warnings the functional dependencies are maintained incrementally. A user study showed that these guidance features do improve usability [7]. However, spreadsheet data cannot be assumed to be complete in practice. It would therefore be interesting to extend this approach to partial information, for example under the *no information* interpretation of null values.

It is also an interesting problem to study the properties of Armstrong data trees for FDs in the context of XML. The results of our article should provide valuable information to learn more about the properties of Armstrong data trees for several classes of XML FDs.

Other directions include the problem of dependency inference [62] or data cleaning [82], but also the investigation of extremal problems [128, 129] in the context of partial relations. Finally, all these problems should also be investigated for other interpretations of null values [105, 110, 106, 107, 108].

Acknowledgements. This research is supported by the Marsden fund council from Government funding, administered by the Royal Society of New Zealand. Sven Hartmann is supported by a research grant of the Alfried Krupp von Bohlen and Halbach foundation, administered by the German Scholars organisation.

References

- Codd, E.F.: A relational model of data for large shared data banks. Commun. ACM 13(6) (1970) 377–387
- [2] Abiteboul, S., Hull, R., Vianu, V.: Foundations of Databases. Addison-Wesley (1995)
- [3] Date, C.: Database in depth: relational theory for practitioners. 1 edn. O'Reilly Media (2005)
- [4] Elmasri, R., Navathe, S.: Fundamentals of Database Systems. 6 edn. Addison-Wesley (2010)
- [5] Delobel, C., Adiba, M.: Relational database systems. North Holland (1985)
- [6] Calvanese, D., De Giacomo, G., Lenzerini, M.: Identification constraints and functional dependencies in description logics. In: Proceedings of the Seventeenth International Joint Conference on Artificial Intelligence (IJCAI). (2001) 155–160
- [7] Qian, L., LeFevre, K., Jagadish, H.: CRIUS: User-friendly database design. Proceedings of the VLDB Endowment 4(2) (2011) 81–92
- [8] Rosenthal, A., Reiner, D.: Tools and transformations rigorous and otherwise for practical database design. ACM Trans. Database Syst. 19(2) (1994) 167–211

- [9] Thalheim, B.: Entity-Relationship modeling. Springer, Heidelberg, Germany (2000)
- [10] Toman, D., Weddell, G.: On keys and functional dependencies as first-class citizens in description logics. J. Autom. Reasoning 40(2-3) (2008) 117–132
- [11] Cohen, E.S.: The theory, practice and methodology of relational database design and programming. http://www.openlineconsult.com/ (2008)
- [12] Ling, T.: A normal form for Entity-Relationship diagrams. In: Proceedings of the Fourth International Conference on the Entity-Relationship Approach (ER). (1985) 24–35
- [13] Rauh, O., Stickel, E.: Standard transformations for the normalization of ER schemata. Inf. Syst. 21(2) (1996) 187–208
- [14] Kolahi, S., Libkin, L.: An information-theoretic analysis of worst-case redundancy in database design. ACM Trans. Database Syst. 35(1) (2010)
- [15] Yu, C., Jagadish, H.: XML schema refinement through redundancy detection and normalization. VLDB J. 17(2) (2008) 203–223
- [16] Fan, W., Ma, S., Hu, Y., Liu, J., Wu, Y.: Propagating functional dependencies with conditions. PVLDB 1(1) (2008) 391–407
- [17] Klug, A., Price, R.: Determining view dependencies using tableaux. ACM Trans. Database Syst. 7(3) (1982) 361–380
- [18] Tan, H., Zhao, Y.: Automated elicitation of functional dependencies from source codes of database transactions. Information & Software Technology 46(2) (2004) 109–117
- [19] Fan, W., Geerts, F., X., J.: A revival of integrity constraints for data cleaning. PVLDB 1(2) (2008) 1522–1523
- [20] Fan, W., Geerts, F., X., J.: Semandaq: a data quality system based on conditional functional dependencies. PVLDB 1(2) (2008) 1460–1463
- [21] Marnette, B., Mecca, G., Papotti, P.: Scalable data exchange with functional dependencies. PVLDB 3(1) (2010) 105–116
- [22] Cali, A., Calvanese, D., De Giacomo, G., Lenzerini, M.: Data integration under integrity constraints. Inf. Syst. 29(2) (2004) 147–163
- [23] Miller, R., Hernandez, M., Haas, L., Yan, L.L., Ho, C., Fagin, R., Popa, L.: The Clio project: Managing heterogeneity. SIGMOD Record 30(1) (2001) 78–83
- [24] Wang, D., Dong, X., Sarma, A., Franklin, M., Halevy, A.: Functional dependency generation and applications in pay-as-you-go data integration systems. In: Proceedings of the 12th International Workshop on the Web and Databases (WebDB), ACM (2009)

- [25] Bisbal, J., Grimson, J.: Database sampling with functional dependencies. Information & Software Technology 43(10) (2001) 607–615
- [26] De Marchi, F., Petit, J.M.: Semantic sampling of existing databases through informative Armstrong databases. Inf. Syst. 32(3) (2007) 446–457
- [27] Lechtenbörger, J., Vossen, G.: Multidimensional normal forms for data warehouse design. Inf. Syst. 28(5) (2003) 415–434
- [28] Kimura, H., Huo, G., Rasin, A., Madden, S., Zdonik, S.: Correlation maps: A compressed access method for exploiting soft functional dependencies. PVLDB 2(1) (2009) 1222–1233
- [29] Kim, M.S., Whang, K.Y., Lee, J.G., Lee, M.J.: Structural optimization of a fulltext n-gram index using relational normalization. VLDB J. 17(6) (2008) 1485–1507
- [30] Koch, C., Olteanu, D.: Conditioning probabilistic databases. PVLDB 1(1) (2008) 313–325
- [31] Wolf, G., Kalavagattu, A., Khatri, H., Balakrishnan, R., Chokshi, B., Fan, J., Chen, Y., Kambhampati, S.: Query processing over incomplete autonomous databases: query rewriting using learned data dependencies. VLDB J. 18(5) (2009) 1167–1190
- [32] Beskales, G., Ilyas, I., Golab, L.: Sampling the repairs of functional dependency violations under hard constraints. PVLDB 3(1) (2010) 197–207
- [33] Molinaro, C., Greco, S.: Polynomial time queries over inconsistent databases with functional dependencies and foreign keys. Data Knowl. Eng. 69(7) (2010) 709–722
- [34] Cheng, Q., Gryz, J., Koo, F., Leung, C., Liu, L., Qian, X., Schiefer, K.: Implementation of two semantic query optimization techniques in DB2 universal database. In: Proceedings of the 25th International Conference on Very Large Data Bases (VLDB). (1999) 687–698
- [35] Deutsch, A., Ludäscher, B., Nash, A.: Rewriting queries using views with access patterns under integrity constraints. Theor. Comput. Sci. **371**(3) (2007) 200–226
- [36] Paulley, G.: Exploiting functional dependence in query optimization. Technical Report UW-CS-2000-11, University of Waterloo, Waterloo, Canada (2000)
- [37] Biskup, J., Embley, D., Lochner, J.: Reducing inference control to access control for normalized database schemas. Inf. Proc. Letters 106(1) (2008) 8–12
- [38] Wang, H., Liu, R.: Privacy-preserving publishing data with full functional dependencies. In: Proceedings of the 15th International Conference on Database Systems for Advanced Applications, Part II (DASFAA). Volume 5982 of Lecture Notes in Computer Science., Springer (2010) 176–183

- [39] Alexe, B., Kolaitis, P., Tan, W.C.: Characterizing schema mappings via data examples. In: Proceedings to the Twenty-Ninth ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems (PODS). (2010) 261–271
- [40] Beeri, C., Dowd, M., Fagin, R., Statman, R.: On the structure of Armstrong relations for functional dependencies. J. ACM 31(1) (1984) 30–46
- [41] De Marchi, F., Lopes, S., Petit, J.M., Toumani, F.: Analysis of existing databases at the logical level: the DBA companion project. SIGMOD Record 32(1) (2003) 47–52
- [42] Mannila, H., Räihä, K.J.: Design by example: An application of Armstrong relations. J. Comput. Syst. Sci. 33(2) (1986) 126–141
- [43] Fagin, R.: Armstrong databases. Technical Report RJ3440(40926), IBM Research Laboratory, San Jose, California, USA (1982)
- [44] Langeveldt, W.D., Link, S.: Empirical evidence for the usefulness of Armstrong relations in the acquisition of meaningful functional dependencies. Inf. Syst. 35(3) (2010) 352–374
- [45] CA Technologies: ERwin Data Modeler methods guide. https://support.ca.com/cadocs/0/e002961e.pdf, page 86 (2011)
- [46] Atzeni, P., Morfuni, N.: Functional dependencies and constraints on null values in database relations. Information and Control 70(1) (1986) 1–31
- [47] Lien, E.: On the equivalence of database models. J. ACM 29(2) (1982) 333–362
- [48] Hartmann, S., Link, S.: When data dependencies over SQL tables meet the Logics of Paradox and S-3. In: Proceedings to the Twenty-Ninth ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems (PODS). (2010) 317–326
- [49] Fagin, R., Vardi, M.: Armstrong databases for functional and inclusion dependencies. Inf. Process. Lett. 16(1) (1983) 13–19
- [50] Armstrong, W.W.: Dependency structures of database relationships. Information Processing 74 (1974) 580–583
- [51] Demetrovics, J.: On the equivalence of candidate keys with Sperner systems. Acta Cybern. 4 (1980) 247–252
- [52] Fagin, R., Vardi, M.: The theory of data dependencies an overview. In: Proceedings of the 11th Colloqium on Automata, Languages and Programming (ICALP). Number 172 in Lecture Notes in Computer Science, Springer (1984) 1–22
- [53] Maier, D.: The Theory of Relational Databases. Computer Science Press (1983)
- [54] Thalheim, B.: Dependencies in relational databases. Teubner (1991)

- [55] Codd, E.F.: Further normalization of the database relational model. In: Proceedings of the Courant Computer Science Symposia 6: Data Base Systems. (1972) 33–64
- [56] Beeri, C., Bernstein, P.: Computational problems related to the design of normal form relational schemas. ACM Trans. Database Syst. 4(1) (1979) 30–59
- [57] Diederich, J., Milton, J.: New methods and fast algorithms for database normalization. ACM Trans. Database Syst. 13(3) (1988) 339–365
- [58] Beeri, C., Mendelzon, A.O., Sagiv, Y., Ullman, J.D.: Equivalence of relational database schemes. SIAM J. Comput. 10(2) (1981) 352–370
- [59] Bernstein, P.: Synthesizing third normal form relations from functional dependencies. ACM Trans. Database Syst. 1(4) (1976) 277–298
- [60] Beeri, C., Bernstein, P.A., Goodman, N.: A sophisticate's introduction to database normalization theory. In: Proceedings of the Fourth International Conference on Very Large Data Bases (VLDB), IEEE Computer Society (1978) 113–124
- [61] Biskup, J., Dayal, U., Bernstein, P.: Synthesizing independent database schemas. In: Proceedings of the International Conference on Management of Data (SIG-MOD). (1979) 143–151
- [62] Mannila, H., Räihä, K.J.: Algorithms for inferring functional dependencies from relations. Data Knowl. Eng. 12(1) (1994) 83–99
- [63] Biskup, J.: Security in computing systems. Springer (2009)
- [64] Deutsch, A., Popa, L., Tannen, V.: Query reformulation with constraints. SIG-MOD Record 35(1) (2006) 65–73
- [65] Arenas, M., Fan, W., Libkin, L.: On the complexity of verifying consistency of XML specifications. SIAM J. Comput. 38(3) (2008) 841–880
- [66] Arenas, M., Libkin, L.: A normal form for XML documents. ACM Trans. Database Syst. 29(1) (2004) 195–232
- [67] Bojanczyk, M., Muscholl, A., Schwentick, T., Segoufin, L.: Two-variable logic on data trees and XML reasoning. J. ACM 56(3) (2009)
- [68] Buneman, P., Davidson, S., Fan, W., Hara, C., Tan, W.: Reasoning about keys for XML. Inf. Syst. 28(8) (2003) 1037–1063
- [69] Fischer, P.C., Saxton, L.V., Thomas, S.J., Van Gucht, D.: Interactions between dependencies and nested relational structures. J. Comput. Syst. Sci. 31(3) (1985) 343–354
- [70] Gottlob, G., Pichler, R., Wei, F.: Tractable database design through bounded treewidth. Inf. Syst. 35(3) (2010) 278–298

- [71] Hartmann, S., Link, S.: Characterising nested database dependencies by fragments of propositional logic. Ann. Pure Appl. Logic 152(1-3) (2008) 84–106
- [72] Hartmann, S., Link, S.: Efficient reasoning about a robust XML key fragment. ACM Trans. Database Syst. 34(2) (2009)
- [73] Hartmann, S., Link, S.: Numerical constraints on XML data. Inf. Comput. 208(5) (2010) 521–544
- [74] Jensen, C., Snodgrass, R., Soo, M.: Extending existing dependency theory to temporal databases. IEEE Trans. Knowl. Data Eng. 8(4) (1996) 563–582
- [75] Kolahi, S.: Dependency-preserving normalization of relational and XML data. J. Comput. Syst. Sci. 73(4) (2007) 636–647
- [76] Levene, M., Loizou, G.: Axiomatisation of functional dependencies in incomplete relations. Theor. Comput. Sci. 206(1-2) (1998) 283–300
- [77] Sözat, M., Yazici, A.: A complete axiomatization for fuzzy functional and multivalued dependencies in fuzzy database relations. ACM Fuzzy Sets and Systems 117(2) (2001) 161–181
- [78] Tari, Z., Stokes, J., Spaccapietra, S.: Object normal forms and dependency constraints for object-oriented schemata. ACM Trans. Database Syst. 22 (1997) 513– 569
- [79] Vincent, M., Liu, J., Liu, C.: Strong functional dependencies and their application to normal forms in XML. ACM Trans. Database Syst. 29(3) (2004) 445–462
- [80] Weddell, G.: Reasoning about functional dependencies generalized for semantic data models. ACM Trans. Database Syst. 17(1) (1992) 32–64
- [81] Wijsen, J.: Temporal FDs on complex objects. ACM Trans. Database Syst. 24(1) (1999) 127–176
- [82] Fan, W., Geerts, F., Jia, X., Kementsietsidis, A.: Conditional functional dependencies for capturing data inconsistencies. ACM Trans. Database Syst. 33(2) (2008)
- [83] Davidson, S., Fan, W., Hara, C.: Propagating XML constraints to relations. J. Comput. Syst. Sci. 73(3) (2007) 316–361
- [84] Chomicki, J.: Consistent query answering: Five easy pieces. In: Proceedings of the 11th International Conference on Database Theory. Number 4353 in Lecture Notes in Computer Science, Springer (2007) 1–17
- [85] Fagin, R., Kolaitis, P., Popa, L., Tan, W.: Reverse data exchange: coping with nulls. In: Proceedings of the Twenty-Eight ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems (PODS). (2009) 23–32

- [86] Libkin, L.: Data exchange and incomplete information. In: Proceedings of the Twenty-Fifth ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems (PODS). (2006) 60–69
- [87] Miklau, G., Suciu, D.: A formal analysis of information disclosure in data exchange. J. Comput. Syst. Sci. 73(3) (2007) 507–534
- [88] Katona, G., Tichler, K.: Some contributions to the minimum representation problem of key systems. In: Proceedings of the 4th International Symposium on Foundations of Information and Knowledge Systems (FoIKS). Volume 3861 of Lecture Notes in Computer Science., Springer (2006) 240–257
- [89] Demetrovics, J., Thi, V.: Some remarks on generating Armstrong and inferring functional dependencies relation. Acta Cybern. 12(2) (1995) 167–180
- [90] Fagin, R.: Horn clauses and database dependencies. J. ACM 29(4) (1982) 952–985
- [91] Demetrovics, J., Thi, V.: Armstrong relations, functional dependencies and strong dependencies. Computers and Artificial Intelligence 14(3) (1995)
- [92] Gottlob, G., Libkin, L.: Investigation on Armstrong relations, dependency inference, and excluded functional dependencies. Acta Cybern. 9(4) (1990) 385–402
- [93] Silva, A., Melkanoff, M.: A method for helping discover the dependencies of a relation. In: Advances in Data Base Theory. (1979) 115–133
- [94] Enders, A., Romback, H.: A Handbook of Software and Systems Engineering; Empirical Overservations, Laws and Theories. Addison-Wesley (2003)
- [95] Martin, J.: Information Engineering. Prentice Hall (1989)
- [96] Standish Group: Unfinished voyages. The Standish Group International, available on-line at http://www.standishgroup.com/sample_research/ unfinsihed_voyages_1.php (1995)
- [97] Boehm, B.: Software Engineering Economics. Prentice Hall (1981)
- [98] Zultner, R.: The deming way: Total quality management for software. In: Proceedings of Total Quality Management for Software. (1992) 134–145
- [99] Codd, E.F.: Extending the database relational model to capture more meaning. ACM Trans. Database Syst. 4(4) (1979) 397–434
- [100] Imielinski, T., Lipski Jr., W.: Incomplete information in relational databases. J. ACM 31(4) (1984) 761–791
- [101] Levene, M., Loizou, G.: Database design for incomplete relations. ACM Trans. Database Syst. 24(1) (1999) 80–125

- [102] Imielinski, T., Van der Meyden, R., Vadaparty, K.: Complexity tailored design: a new design methodology for databases with incomplete information. J. Comput. Syst. Sciences 51(3) (1995) 405–432
- [103] Libkin, L., Wong, L.: Semantic representations and query languages for Or-sets. J. Comput. Syst. Sciences 52(1) (1996) 125–142
- [104] Vadaparty, K., Naqvi, S.: Using constraints for efficient query processing in nondeterministic databases. IEEE Trans. Knowl. Data Eng. 7(6) (1995) 850–864
- [105] Codd, E.F.: Understanding relations. ACM SIGFIDET FDT Bulletin 7(3-4) (1975) 23–28
- [106] Grahne, G.: Dependency satisfaction in databases with incomplete information. In: Proceedings of the Tenth International Conference on Very Large Data Bases (VLDB), Morgan Kaufmann (1984) 37–45
- [107] Grant, J.: Null values in a relational data base. Inf. Process. Lett. $\mathbf{6}(5)$ (1977) 156–157
- [108] Makinouchi, A.: A consideration on normal form of not-necessarily-normalised relation in the relational data model. In: Proceedings of the Third International Conference on Very Large Data Bases (VLDB). (1977) 447–453
- [109] Zaniolo, C.: Database relations with null values. J. Comput. Syst. Sci. 28(1) (1984) 142–166
- [110] Gottlob, G., Zicari, R.: Closed world databases opened through null values. In: Proceedings of the Fourteenth International Conference on Very Large Data Bases (VLDB), Morgan Kaufmann (1988) 50–61
- [111] Schaerf, M., Cadoli, M.: Tractable reasoning via approximation. Artif. Intell. 74 (1995) 249–310
- [112] Link, S.: On the implication of multivalued dependencies in partial database relations. Int. J. Found. Comput. Sci. 19(3) (2008) 691–715
- [113] Ginsburg, S., Hull, R.: Characterizations for functional dependency and Boyce-Codd normal form families. Theor. Comput. Sci. 26 (1983) 243–286
- [114] Lucchesi, C., Osborn, S.: Candidate keys for relations. J. Comput. Syst. Sci. 17(2) (1978) 270–279
- [115] Maier, D.: Minimum covers in relational database model. J. ACM $\mathbf{27}(4)$ (1980) 664–674
- [116] Bernstein, P.A., Goodman, N.: What does Boyce-Codd Normal Form do? In: Proceedings of the Sixth International Conference on Very Large Data Bases (VLDB). (1980) 245–259

- [117] Vincent, M.: Semantic foundations of 4NF in relational database design. Acta Inf. 36(3) (1999) 173–213
- [118] Bhargava, G., Goel, P., Iyer, B.: Hypergraph-based reorderings of outer join queries with complex predicates. In: Proceedings of the International Conference on Management of Data (SIGMOD), ACM (1995) 304–315
- [119] Hasan, W., Pirahesh, H.: Query rewrite optimization in STARBURST. Technical Report RJ6367, IBM Research Laboratory, San Jose, California, USA (1988)
- [120] Ilyas, I., Markl, V., Haas, P., Brown, P., Aboulnaga, A.: CORDS: Automatic discovery of correlations and soft functional dependencies. In: Proceedings of the International Conference on Management of Data (SIGMOD), ACM (2004) 647– 658
- [121] Paulley, G., Larson, P.A.: Exploiting uniqueness in query optimization. In: Proceedings of the Tenth International Conference on Data Engineering (ICDE), IEEE Computer Society (1994) 68–79
- [122] Simmen, D., Shekita, E., Malkemus, T.: Fundamental techniques for order optimization. In: Proceedings of the International Conference on Management of Data (SIGMOD). (1996) 57–67
- [123] Biskup, J., Weibert, T.: Keeping secrets in incomplete databases. Int. J. Inf. Sec. 7(3) (2008) 199–217
- [124] Farkas, C., Jajodia, S.: The inference problem: a survey. SIGKDD Explorations 4(2) (2002) 6–11
- [125] Levene, M., Loizou, G.: Null inclusion dependencies in relational databases. Inf. Comput. 136(2) (1997) 67–108
- [126] Mannila, H., Räihä, K.J.: Design of Relational Databases. Addison-Wesley (1992)
- [127] Fagin, R.: The decomposition versus the synthetic approach to relational database design. In: Proceedings of the Third International Conference on Very Large Data Bases (VLDB). (1977) 441–446
- [128] Demetrovics, J., Katona, G., Miklos, D., Thalheim, B.: On the number of independent functional dependencies. In: Proceedings of the 4th International Symposium on Foundations of Information and Knowledge Systems (FoIKS). Number 3861 in Lecture Notes in Computer Science, Springer (2006) 83–91
- [129] Engel, K.: Sperner theory. Cambridge Univ. Press (1997)